



# Handelshøyskolen BI

## GRA 19703 Master Thesis

Thesis Master of Science 100% - W

### Predefinert informasjon

<b>Startdato:</b>	09-01-2023 09:00 CET	<b>Termin:</b>	202310
<b>Sluttdato:</b>	03-07-2023 12:00 CEST	<b>Vurderingsform:</b>	Norsk 6-trinns skala (A-F)
<b>Eksamensform:</b>	T		
<b>Flowkode:</b>	202310  11184  IN00  W  T		
<b>Intern sensor:</b>	(Anonymisert)		

### Deltaker

Navn: Joakim Gjølstad og Nikolai August Stellander Skants

### Informasjon fra deltaker

Tittel \*: Artificial Intelligence in Organizational Decision-Making

Navn på veileder \*: Vegard Kolbjørnsrud

Inneholder besvarelsen konfidensielt materiale?  Nei  Ja

Kan besvarelsen offentliggjøres?  Ja  Nei

### Gruppe

Gruppenavn: (Anonymisert)

Gruppenummer: 26

Andre medlemmer i gruppen:

BI Norwegian Business School  
Oslo, Spring 2023

# Artificial Intelligence in Organizational Decision-Making

**Joakim Gjølstad**

**Nickolai August Steller Skants**

**Campus:**

BI Oslo

**Examination Code and Name:**

GRA 19703 – Master Thesis

**Programme:**

Master of Science in Business - Major in Strategy

**Supervisor:**

Vegard Kolbjørnsrud

## Acknowledgements

We would like to express our sincerest gratitude to our academic supervisor, Dr. Vegard Kolbjørnsrud for his continuous guidance, support, and insightful critiques throughout the course of this master's thesis. His expertise and dedication have been invaluable in the pursuit of this study. The knowledge and wisdom he shared with us have not only shaped this thesis but also enhanced our understanding of the intersection between business and technology.

We would also give a special thanks to all the interviewees who willingly and generously participated in this study. Their willingness to participate and share their personal insights was crucial to the development of this research. These contributions have made this thesis not only possible but also more enriched and nuanced than we ever could have dreamt of.

Finally, a special mention to our friends and family for their continuous support and encouragement during the last year. Their belief in our abilities was a constant source of motivation.

Thank you all for contributing to the completion of this journey. The lessons learned and experiences gathered along the way will forever resonate in both our professional and personal lives.

## Executive summary

Artificial Intelligence (AI) has been around since the 1950s, but its potential is yet to be fully realized. Our findings suggests that AI can augment human intelligence and allow for more efficient and rational decision-making processes, given that the contingencies are effectively addressed. Inspired by the literature on organizational decision-making and AI, our study is grounded in the belief that understanding AI's impact on decision-making requires not just a technical understanding of AI, but also a deep exploration of its integration and implications within an organizational context. As such, we conducted interviews with various experts and users, utilizing a cross-case analysis to answer the following research question:

***How do machine learning and natural language processing augment operational decision-making processes in organizations?***

This gave us valuable insights into the perceived definition of AI, attitudes and expectations toward AI, and the benefits and challenges of AI in organizational decision-making. Juxtaposing this with the theoretical foundation, we discuss our most important theoretical findings: 1. Rationality and Accuracy, 2. Trust in AI, 3. Organizational Structure and Strategic Goals, and 4. Problem Comprehension. Following this, we provide a list of managerial recommendations, in addition to addressing possible limitations with the study and suggestions for further research. As such, we provide a holistic understanding of the dynamic factors and intricacies central to successfully augmenting organizational decision-making processes with AI.

## Table of Contents

<b>1.0 Introduction</b>	<b>5</b>
1.1 Purpose and Objectives	6
1.2 Research Question	7
1.3 Structure	9
<b>2.0 Literature Review</b>	<b>9</b>
2.1 Artificial Intelligence	10
2.1.1 Definitions, Historical Evolution, and Recent Developments	10
2.1.2 AI fundamentals	11
2.1.2.1 Predictive Analytics and Classification Analytics	13
2.1.2.2 Natural Language Processing	14
2.2 Organizational Decision-Making	15
2.2.1 Bounded Rationality and Heuristics	17
2.2.2 Organizational and Situational Context	21
2.3 Artificial Intelligence in Organizational Decision-Making	22
2.3.1 Complementarity Between Humans and Machines	23
2.3.2 Decision-Making Structures	24
2.3.3 Algorithmic Bias, Transparency, and Data Issues	26
2.3.4 Organizational Trust in AI	28
2.3.5 Successfully Implementing AI	29
2.3.6 The Existing State of Literature	31
<b>3.0 Methodology</b>	<b>34</b>
3.1 Research Design	34
3.1.1 Qualitative Research Design	34
3.1.2 Case Sampling	36
3.1.3 Case Descriptions	38
3.1.3.1 Cases with Predictive Analytics	39
3.1.3.2 Cases with Natural Language Processing	41
3.1.3.3 Cases with Classification Analytics	42
3.2 Data Collection	43
3.3 Analytical Process	44
3.4 Ensuring Quality in the Study	47

3.5 Ethical Considerations.....	48
<b>4.0 Findings and analysis.....</b>	<b>49</b>
4.1 The Perceived Definition of AI .....	50
4.2 Attitudes and Expectations toward AI.....	51
4.3 The Benefits of AI in Organizational Decision-Making .....	53
4.3.1 Accuracy and Efficiency .....	53
4.3.2 Expanded Information Base and Sustainability .....	56
4.3.3 Increased Adaptability, Reliability and Preservation of Knowledge ....	58
4.4 The Challenges of AI in Organizational Decision-Making.....	60
4.4.1 Problem Identification and Comprehension.....	60
4.4.2 Trust in AI .....	61
4.4.3 Costs and Complexity of Data Management.....	63
4.4.4 Regulatory and Economic Uncertainty .....	65
4.4.5 Organizational Complexity and Inertial Forces .....	67
<b>5.0 Discussion.....</b>	<b>69</b>
5.1 Theoretical Implications .....	69
5.1.1 Rationality and Accuracy .....	69
5.1.2 Trust in AI .....	71
5.1.2.1 Managing Trust in Organizations .....	72
5.1.2.2 Attitudes and Knowledge Towards AI .....	74
5.1.3 Organizational Structure and Strategic Goals .....	75
5.1.4 Problem Comprehension .....	77
5.2 Managerial Implications .....	78
5.3 Limitations .....	80
5.4 Suggestions for Further Research .....	81
<b>6.0 Conclusion .....</b>	<b>83</b>
<b>7.0 References .....</b>	<b>85</b>
<b>8.0 Appendix: A Semi-Structured Interview-Guide .....</b>	<b>99</b>

# 1.0 Introduction

Decision-making processes can be defined as the process of selection between alternatives, in which the choice implies a commitment to action (Langley et al., 1995). Only 50-75 years ago, the central processor of decision-making in businesses was human judgment based on intuition, experience, and gut instinct (Colson, 2019). However, thanks to multiple decades of research on decision-making, we now have a detailed picture of how human judgment is bounded and that intuition-based decision-making is inherently flawed (eg. Milkman et al., 2009; Simon, 1979; Tversky & Kahneman, 1974; 1986). These flaws stem from cognitive biases and judgmental heuristics common to human decision-making under uncertainty, such as the ones linked to representativeness, availability, adjustment, and anchoring, to name a few (Tversky & Kahneman, 1974). Although they can lead to systematic and predictable errors (eg. Tversky & Kahneman, 1974; Simon, 1957), they can allow for efficient and accurate decision-making processes in several contexts (Czerlinski, 1999; Gigerenzer & Gaissmaier, 2011).

One way of addressing the fallacies of human decision-making is to utilize Artificial Intelligence, as it is less prone to cognitive biases of humans (Colson, 2019). Kolbjørnsrud et al. (2017) define Artificial intelligence as “IT systems that sense, comprehend, act and learn” (p. 43). Its potential lies in the fact that, compared to humans, AI-based machines are faster, more accurate, consistent, and do not get tired (Shrestha et al., 2019). With the recent breakthroughs in machine learning techniques, particularly in deep learning (LeCun et al., 2015), the global market is expected to be valued at almost \$1400 billion US dollars by 2029, growing over elevenfold from today’s value (Fortune Business Insight, 2022). Incorporating AI algorithms into workflows can allow for more consistent and objective decisions in certain contexts, but its utility might be limited due to bias, data, and transparency issues (Colson, 2019; Shrestha et al., 2019). In addition, these systems also lack intuition, emotion, and cultural sensitivity (Cremer & Kasparov, 2021). As a result, AI has the capability, when used correctly, to augment human intelligence. By ‘augmenting human intelligence’, we build on

the work of Engelbart (1962), that augmenting human intelligence means to “increase the capability of a man to approach a complex problem situation... and to derive solutions to problems” (p. 1). As such, based on existing research, combining AI and human intelligence seems to be the golden standard for the years to come.

With the introduction of ground-breaking natural language processing chatbots such as ChatGPT (OpenAI, 2023a), the impact of AI is inevitable. The rapid development of advanced and opaque AI solutions has, however, stirred up noise amongst academics, companies, and nations, evident by Italy’s recent ban on the use of ChatGPT (Financial Express, 2023; Jawad, 2023; McCallum, 2023). Although the ban was lifted rather quickly, it highlights the tension around the use of AI (Robertson, 2023), as it is being developed faster than regulatory organs can follow. This has led to an increasingly popular movement that aims to pause any development of models more powerful than GPT-4 for six months (Future of Life Institute, 2023; Kahn, 2023). With more than 33 000 signatures (at the time of writing) (Future of Life Institute, 2023), including major names such as Elon Musk and Steve Wozniak (Kahn, 2023), the proposed ban aims to allow for careful consideration of potential risks of advanced AI models, in addition to allowing for the development of safety measures (Future of Life Institute, 2023).

## 1.1 Purpose and Objectives

The aim of this dissertation is to analyze several AI use cases across different companies and stakeholders, based on AI and decision-making literature. More specifically, we want to get a thorough understanding of the organizational advantages and disadvantages that promote or inhibit the development and the implementation of AI. As about 80-90% of all AI investments yield no organizational value (Gartner, 2018), a mission of ours is to contribute to the rate’s decline. This includes, but is not limited to, examining how the AI innovation took place in organizations; the change management aspects of introducing AI; problem-solving aspects - to observe whether AI was applied with a well-reasoned problem in mind; and what advice the different case participants



have for other companies that plan or want to implement AI into their operations. Therefore, our study is not only descriptive, but also highly prescriptive.

As all organizations are different (Barney, 1991) and thus have different prerequisites for reaping the benefits of AI and handling the disadvantages of arising challenges, we expect that there will be differences as to how AI will shape companies and their competitive environments. Although the prevalence of AI has increased drastically in the last couple of years, our understanding of AI's impact on companies is limited (Scarpetta, 2023). Therefore, our study aims to enhance the understanding of the intricacies of AI and how it affects organizations by filling the gaps in the existing literature, as identified in subchapter 2.3.6.

## 1.2 Research Question

University of Southern California (n.d.) defines a research question as “a definite or clear expression [statement] about an area of concern, a condition to be improved upon... that exists in scholarly literature, in theory, or within existing practice that points to a need for meaningful understanding and deliberate investigation.” Furthermore, this provides the basis for the choice of research design and procedures. To illuminate the purpose and objective of the study, we arrived at the following research question:

*How do machine learning and natural language processing augment operational decision-making processes in organizations?*

Although we acknowledge the fact that machine learning is a central part of natural language processing (NLP) models, in this thesis we divide them into two different areas to easier clarify and separate the different usage areas in our cases. The reason is that these are fundamentally different, as machine learning encompasses a broader range of techniques and applications beyond language understanding, while NLP focuses specifically on the challenges posed by human language and requires specialized methods to extract meaning and context.

As of today, the use of AI is mostly limited to operational decision-making, with strategic decision-making lagging behind, but rapidly growing (Borges et al., 2021; Perifanis & Kitsios, 2023). As all the cases we found and analyzed are operational use cases of AI, this became a natural limitation of our study. In addition, we limited our study to machine learning (ML) algorithms and natural language processing (NLP), as these were the prevalent AI technologies in our chosen cases.

Despite our main concern with how AI *augments* decision-making, we recognize that there are several hindrances and limitations to successful implementation. These are critical to understand as managers and employees in every organization face the difficulties of implementing AI in decision-making to avoid being at a competitive disadvantage in the years to come. These difficulties can include, but are not limited to, biased algorithms, black box and explainability challenges, and other technical, ethical, economic, and regulatory constraints (Snow & Fjeldstad, 2023; Shrestha et al., 2019). As these factors are critical to get right for AI to yield value for companies and users, they became a natural part of our study.

In understanding the impact of AI on organizational decision-making, it is essential to distinguish between augmentation, as we build our thesis on, and automation (Davenport & Kirby, 2016). Automation refers to the process in which tasks, primarily those considered to be routine-based, predictable, and less complex, are entirely transferred from human workers to AI systems, thereby reducing human intervention to a bare minimum (Brynjolfsson & McAfee, 2014). Augmentation, on the other hand, can be a more collaborative model, in which humans utilize AI tools to enhance their cognitive abilities and productivity by complementing each other's strengths and weaknesses (Davenport & Ronanki, 2018; Engelbart, 1962). In essence, one is then able to utilize the best abilities from both worlds, and as such, enhance an organization's decision-making capabilities (Agrawal, Gans, & Goldfarb, 2018; Davenport & Ronanki, 2018; Engelbart, 1962).

While the formulation of our research question is in the present tense, we also want to address how augmentation will work in the years to come as sustainability has arisen as a critical global challenge. As sustainability related issues are addressed most effectively through changes in attitudes or technological innovation (Xiao & Su, 2022), we aim to explore how organizations view AI tools as an approach to improve sustainability. At the same time, we aspire to provide additional understandings of the potential trade-offs and unintended negative consequences of AI implementation on sustainability.

### 1.3 Structure

Our paper begins by reviewing the current literature on AI and organizational decision-making independently, and then connecting these in the subsequent section to establish a thorough understanding. Following this, we explain the research methodology and the decisions we have made while designing and conducting this study. This includes an overview of the 11 selected cases and more detailed descriptions of each. Then, we present the most central findings through thematic analysis, before engaging in a discussion where we juxtapose this with the existing literature. We also discuss the limitations of our work and provide suggestions for future research. Finally, we conclude our paper by reviewing our main contributions to the research question.

## 2.0 Literature Review

In our study we aim to observe how machine learning and natural language processing augment operational decision-making processes in organizations. As such, we will review the two streams of literature that build the foundation of this thesis, namely artificial intelligence, and organizational decision-making. Then, we will look at the two separate streams in conjunction in order to evaluate the current academic state of the use of AI in organizational decision-making and its limitations.

## 2.1 Artificial Intelligence

### 2.1.1 Definitions, Historical Evolution, and Recent Developments

Humans can use available information as well as reasoning to solve problems, so why cannot machines do the same thing? This was the question contemplated by Alan Turing (1950) in his seminal paper: "Computing Machinery and Intelligence". It was not until the Dartmouth Conference in 1956, however, that the term AI surfaced for the first time (McCarthy et al., 1955). Since then, the field of AI has experienced periods of alternating progress and stagnation, commonly referred to as "AI Winters" and "AI Summers" (Haenlein & Kaplan, 2019; Russell & Norvig, 2009). In combination with these 'seasons', each era in the history of AI has been characterized by a shift in focus, primarily driven by the constraints of previous methodologies and the advancements in computational resources and data availability (Anyoha, 2020).

Initially, the pioneers of AI aimed to develop machines capable of simulating all aspects of human intelligence by precisely describing any feature of intelligence that a machine could emulate (McCarthy et al., 1955). The ambition to mimic human intelligence, however, proved to be overly ambitious and led to a shift of focus to the development of a rule-based system where predefined rules could be applied to a dataset (Haenlein & Kaplan, 2019). Building on the limitations of this method, machine learning and statistical based methods capable of processing larger amounts of data emerged as a dominant approach by the 1980s and 90s (Nilsson, 2010). In recent years, the field of AI has seen a renewed surge of interest, largely driven by advancements in computational power, storage, and cloud technologies (Jordan & Mitchell, 2015). The biggest factor for the revival, however, is the breakthroughs in machine learning techniques, particularly deep learning (LeCun et al., 2015). By leveraging neural networks with several layers, deep learning algorithms have achieved unprecedented success in areas like image and speech recognition, and natural language processing (LeCun et al., 2015).

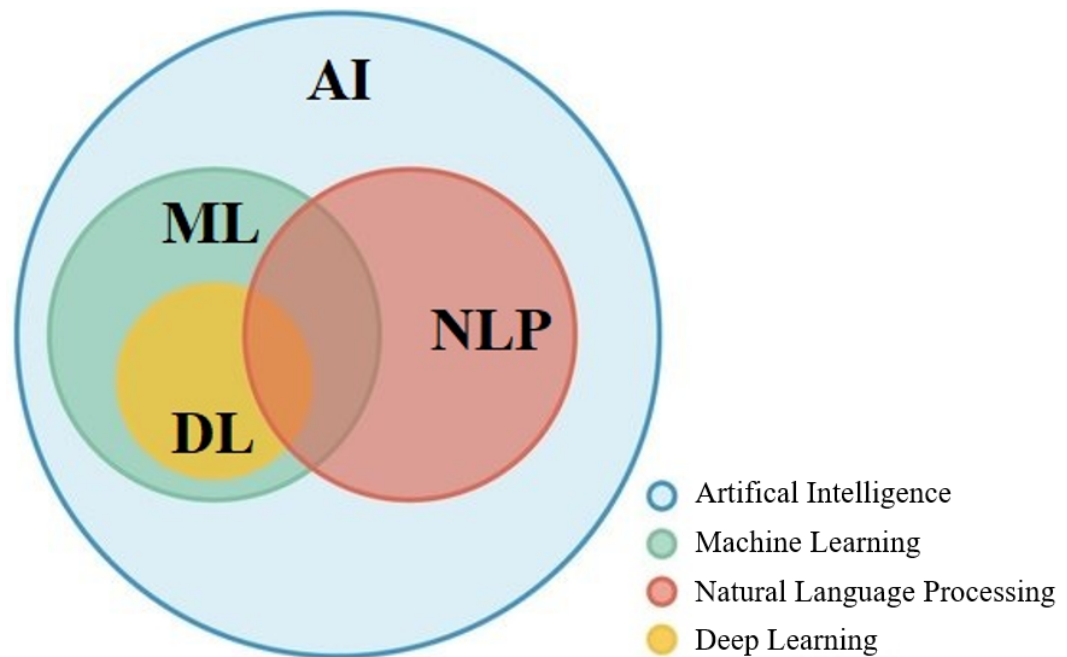
Despite its origins in computer science, AI has evolved to intersect with a variety of disciplines, such as linguistics, psychology, and philosophy (Dartnall, 1994).

This interdisciplinary nature has led to the emergence of multiple AI definitions, where a single definition seems to be unagreeable among scholars, industries, and firms (Littman et al., 2021). Some scholars suggest that the absence of a clear, universal definition has been advantageous for the development of the field since practitioners, researchers, and developers are guided by a general understanding rather than following a strict and limited route (Littman et al., 2021). In this thesis, we build our study on Kolbjørnsrud et al. (2017) definition that "artificial intelligence refers to IT systems that sense, comprehend, act and learn" (p. 43).

Much like the lack of a universally accepted definition of AI, the same phenomenon is present when defining the 'intelligence' component of AI. The spectrum of definitions is broad, from ones leaning heavily on mathematical terminology as discussed by Legg and Hutter (2007), to those that reference the cognitive abilities of animals like cats and dogs, as suggested by Goertzel and Wang (2007). In this study, however, we align with Wang's (1995) definition, that intelligence is "the ability of an information processing system to adapt to its environment with insufficient knowledge and resources" (Wang, 1995, p. 22).

### 2.1.2 AI Fundamentals

Artificial Intelligence can be divided into two main categories: 1) General or strong AI, and 2) Narrow or weak AI. In category 1, machines can replicate human abilities such as understanding, thinking, and feeling. In category 2, machines can perform certain specific tasks with the same, or higher level of proficiency than humans (Fieseler & Bucher, 2022). While General AI has yet to be implemented, many experts believe it will be achieved in the near future (Roser, 2023). According to Russell (2023), there are currently ideas for general-purpose AI that can perform a variety of tasks across different domains, and if humans were to succeed, this can "lift the living standard of everyone on earth".



*Figure 1 - The Relationship Between AI, NLP, ML and DL (Source: AthenaTech, 2019)*

Machine learning is the most popular subset of AI. It is a set of computer algorithms that can detect complex relationships or patterns from empirical data, and among other things and what we are looking into in this thesis, provide the foundation to make reliable decisions (Sarker, 2021a). In the first stage of machine learning development, training, an algorithm is generated from a data set containing inputs and their corresponding outputs. This algorithm can be used to reproduce the examples provided. In the second stage, inference, the algorithm is employed in the application, and user inputs are turned into outputs via the application (Fieseler & Bucher, 2022). Machine learning has three subsets: Supervised learning in which each input example is marked with its expected output value, unsupervised learning where only inputs are given without any correct answer, and reinforcement learning in which an agent learns by gaining feedback from an environment through supervised signals (Fieseler & Bucher, 2022).

Deep learning is a subset of machine learning that has gained a lot of attention in the field of AI recently (Fieseler & Bucher, 2022). It eliminates some of the pre-

processing of inputs by learning feature representations from raw data and is capable of handling large data sets (MIT, 2019). The goal of this method is to reproduce the human brain and its billions of neurons and connected synapses (IBM Cloud Education, 2020). This version works by having inputs pass through multiple hidden layers that calculate weighted sums and biases, triggering neurons in the next layer if threshold values are met according to the activation function (Janiesch et al., 2021). These networks of neurons are trained by “iteratively running a network on examples sampled from very large datasets and then updating the network parameters... to improve performance.” (Fieseler & Bucher, 2022, p. 139). Deep learning has been successful in designing features for various AI applications such as speech recognition, image recognition, and natural language understanding; surpassing many traditional machine learning algorithms (LeCun et al., 2015).

Natural language processing is a subfield of artificial intelligence, that is based on both machine learning and deep learning. NLP aims to enable computers to mimic human linguistic behavior, for it to understand, interpret, and generate human language (Khurana et al., 2022). Since NLP emerged in the 1950s, tech companies and researchers have leveraged various techniques from linguistics, computer science, and artificial intelligence to continuously develop more advanced models capable of deciphering language structures (Nadkarni et al., 2011).

As illustrated in Figure 1, AI consists of a range of different subsets. In this thesis, however, we are mainly focused on predictive analytics, classification analytics and NLP. As such, when mentioning AI in the text, these are the subsets of AI we are referring to.

#### 2.1.2.1 Predictive and Classification Analytics

Predictive analysis and classification analysis are both subfields of machine learning. While prediction analytics enables users to make accurate predictions about future events and trends based on patterns and trends observed in historical data (Kumar & Garg, 2018), classification analytics enables users to close to real-time identify trends and categorize data (Tam et al., 2017). Most commonly, these tools are based on regression and classification models (Lepeniotti et al., 2020).

Rooted in statistics and computer science (Kumar & Garg, 2018), predictive analytics and classification analytics involves several steps (Kelleher et al., 2020). In a simple scenario, the first step typically involves data preparation, whereby the data is cleaned, integrated, and transformed. After this, the data is divided into training and testing sets. The training set is used to develop the model, while the testing set is used to evaluate the model's accuracy. Once the data is ready, it is then important to select the appropriate machine learning algorithm, which is dependent on the problem being addressed. The model can then be developed. This step involves training the selected algorithm on the training set, so it can learn how to make valid predictions. The accuracy of the model can be tested on the testing set. This allows the user to see how well the model has learned underlying patterns in the training set for the purpose of identifying trends. Finally, the model is deployed on a new data set (Kelleher et al., 2020; Kumar & Garg, 2018).

#### 2.1.2.2 Natural Language Processing

The subfield of NLP can be broadly categorized into two main areas: namely natural language understanding and natural language generation (Jurafsky & Martin, 2019). According to Brown et al. (2020), Vaswani et al. (2017) and Wolfe (2022) natural language understanding typically, in a simplified manner, consists of 5 stages, where the first stage is tokenization. This enables the computer to break a document down into separate words. The second step is contextual understanding, where a model interprets the meaning of words and phrases based on their context. While earlier models explicitly used methods such as stemming, lemmatization, speech tagging, entity tagging, newer models such as GPT-3 and GPT-4 learn these abilities from its training data. The third step is self-attention mechanism, which allows the model to weigh the importance of different words in a sentence when trying to predict the next word. The fourth step is training. To train a model to generate text, these steps, together with existing texts and documents are used. According to Delua (2021), one of the most traditionally commonly used methods for training NLP models is supervised learning. Here, a model is trained on a manually labeled dataset by iteratively making predictions on the data and adjusting for the correct answer. Recent advancements in NLP,



however, have been driven by deep learning techniques with semi-supervised and unsupervised learning, such as recurrent neural networks (RNNs) and transformers. The fifth and final step is fine-tuning, which enables the model to train on specific tasks such as summarizing and translation, using a smaller, task-specific dataset.

These models have shown remarkable progress in language understanding and generation tasks. For instance, OpenAI's GPT-3 and GPT-4 models, which are transformer-based language models, have shown impressive results in language generation tasks, such as writing essays, news articles, and poetry (OpenAI, 2023b). In addition to this, it can also answer questions, translate languages, and generate computer code (OpenAI, 2023a).

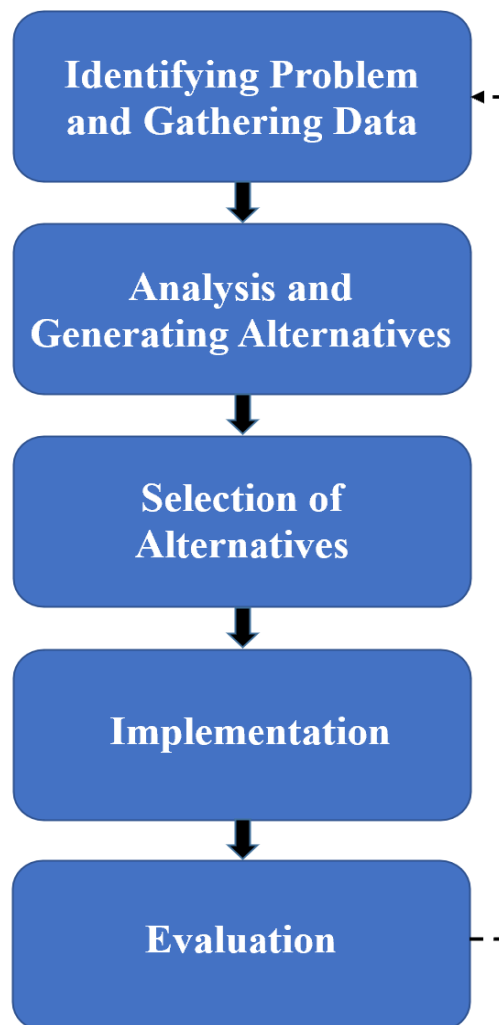
The development of pre-trained language models, such as BERT and RoBERTa, is another notable advancement in the field of language modeling. Models such as these are trained on massive amounts of text data and are therefore able to be fine-tuned to be suitable for specific tasks with relatively little additional data. As illustrated by Devlin et al. (2019) and Khurana et al. (2022), models such as these result in a significant improvement in the performance of NLP tasks such as text classification, sentiment analysis, and named entity recognition.

## 2.2 Organizational Decision-Making

Decision-making processes are integral parts of an organization and are carried out all the time. Although these can be relatively insignificant when viewed separately, they can accumulate and have drastic impacts on organizations (Jacobsen & Thorsvik, 2019). These can encompass smaller and more operational decisions related to how to approach a potential customer, or bigger and more strategic decisions related to what markets to pursue in the future. Getting certain decisions wrong can result in catastrophic consequences for an organization, which is amplified by research that shows that up to half of all decisions made by leaders are incorrect (Nutt, 2008). Moreover, psychological research has highlighted the numerous “limitations” of human decision-making, including aspects of bounded rationality and biases (Boal & Meckler, 2010; March, 1978;

Milkman et al., 2009; Simon, 1979; Tversky & Kahneman, 1974), but also how heuristics can be an efficient decision-making strategy (eg. Czerlinski et al., 1999; Gigerenzer & Gaissmaier, 2011). Other work has shown how decision-making processes interact with organizational context, personal and situational characteristics (eg. Evans, 2008; Jacobsen & Thorsvik, 2019; Lerner et al., 2015; Nutt, 2011), and how earlier decisions can lead to path-dependence (March, 1991; Nelson & Winter, 1982; Sydow et al., 2009).

When it comes to the process of decision-making, we can define it as all actions or assessments that lead to the selection (intentions) and implementation (action) of a decision (Langley, 1995). According to Simon (1960), a decision-making process can be divided into three steps: 1) Intelligence: Identifying a problem and gathering data, 2) Design: Analysis and the generation of alternatives, 3) Choice: Selection of alternatives. Expanding on the work of Simon (1960), we choose to include two additional steps to include the learning aspect of the decision-making process: 4) Implementation of the alternative, and 5) Evaluation of whether the measures worked as they were supposed to (Jacobsen & Thorsvik, 2019). As learning is a cyclical loop rather than a one-directional process, learning loops enable an organization to obtain a deeper understanding by allowing the consolidation of previous knowledge (Dierkes et al., 2003). It is worth noting, however, that not all decision-making processes can be viewed as a loop, which is why Figure 2 includes a dotted line connecting step 5 to step 1.



**Figure 2** – *The Stages of the Decision-Making Processes.* (Adapted from Jacobsen & Thorsvik, 2019, p. 300 and Simon, 1960, p. 1).

### 2.2.1 Bounded Rationality and Heuristics

Common to decision-making theories is the idea that humans act rationally (Simon, 1945). This leads us to the critical distinction between perfect rationality, in which a human is characterized as an “economic man”, and bounded rationality, in which a human is an “administrative man”. The former shows how humans in a predictable and “perfect” world should make decisions, but research has pointed out in various ways that this does not represent real-world scenarios (Boal & Meckler, 2010; Kahneman, 2011; March, 1978; Milkman et al., 2009; Simon, 1964; 1966; 1979; Tversky & Kahneman, 1984; 2000; Weick, 2009). Therefore, bounded rationality was proposed as a more realistic view of human decision-making, which addresses the issues of never having full information

about a case, and that one rarely has clear goals and preferences (Pfiffner, 1960). This model is characterized by *satisficing* decision behavior (Steptoe-Warren et al., 2011, Simon, 1945; 1957), rather than the *maximizing* one, meaning that the alternative that is satisfactory is chosen. This means that the information one is exposed to, in which order alternatives are presented, and the rules and norms that are used to choose between alternatives, are important for what alternative is chosen (Simon, 1947).

	Perfect Rationality	Bounded Rationality
Identifying Problem and Gathering Data	Has a clear understanding of the situation, has clear goals and information about all alternative solutions	Attempts to analyze the situation, determine goals and seeks information about some alternative solutions
Analysis and Generating Alternatives	Considers alternatives according to what the best result is	Considers some alternatives and consequences sequentially
Selection of Alternatives	Selects the alternative that gives the <b>best</b> results in relation to goals	Selects the first alternative that gives a <b>satisfactory</b> result in relation to goals
Implementation	Has extensive knowledge and available resources to implement decision	Lacks knowledge and available resources to implement decision
Evaluation	Evaluates the decision-making process to learn and improve	Neglects evaluation and do not learn and improve as much from the process

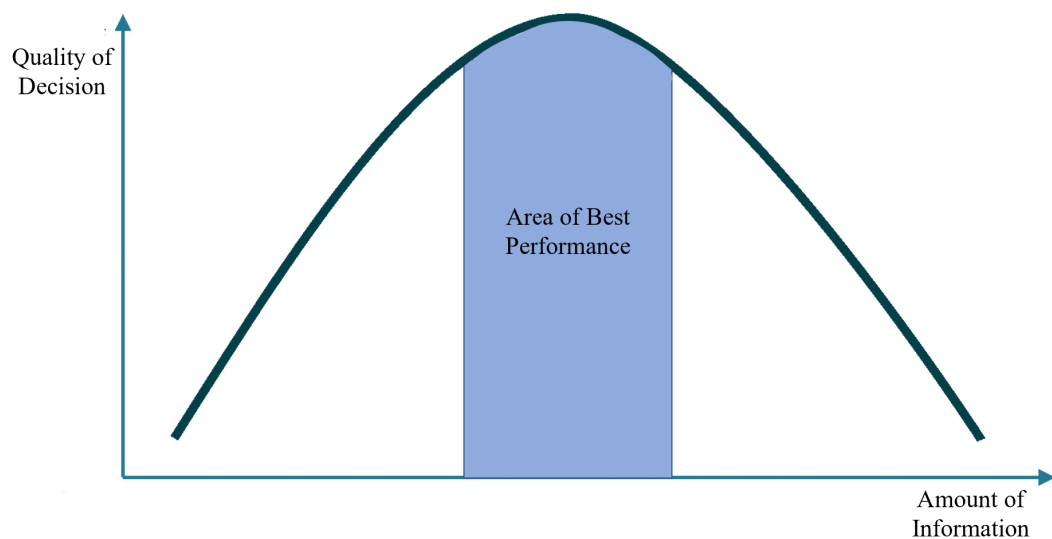
*Table 1: The Key Differences Between Bounded Rationality and Perfect Rationality (Adapted from Jacobsen & Thorsvik, 2019, pp. 302 & 304 and Simon, 1960, p. 1).*

As part of the classical view, heuristics have been linked to errors, lower accuracy, and irrationality, for instance shown in the works of Tversky & Kahneman (1974) on biases and heuristics, and Evans' (2008) work on system 2 theories. First, the former authors argue how the use of heuristics can lead to biases and suboptimal decisions through availability, representativeness, and anchoring and adjustment heuristics. Similarly, the latter author connects heuristics from various theoretical views to errors and biases in system 1 decision-making; "processes that are unconscious, rapid, automatic and high capacity", and connect more accurate and rational decision-making to system 2; "conscious, slow, and deliberative" processes. Common to both is that we would be better off avoiding these biases and errors by striving for perfect rationality or system 2 thinking (Kahneman, 2011).

The classical reason for using heuristics has been that they save effort and time, but at the expense of accuracy (Tversky & Kahneman, 1974; Payne et al., 1993, Shah & Oppenheimer, 2008). This can be rationalized because of the varying importance of decisions, as some cannot justify the effort with system 2 decision-making processes, and due to the cognitive capacity limitations of humans, for instance, that humans can be overloaded with complex information (Tversky & Kahneman, 1974). Although it might seem like humans make countless errors and might be better off handling several aspects of the decision-making process to machines, Kahneman (2011) attributes these errors to the urge to make quick decisions. Further, he proposes that if one neglects this urge and slows down, one will act more rationally.

Heuristics are a central part of bounded rationality and organizational decision-making because the contingencies for rational models are rarely met in the real world (Gigerenzer & Gaissmaier, 2011). Gigerenzer and Gaissmaier (2011) defined these as "strategies that ignore information to make decisions faster, more

frugally, and/or more accurately than complex methods.” (p. 45). The authors stress the importance of these satisficing strategies and why they should be considered as an equally important mental tool alongside logic and statistics. The reason for this is that “heuristics can be more accurate than more complex strategies even though they process less information”, characterized as “less-is-more” effects (p. 474). As illustrated in Figure 3, an inverse U-shaped relationship was found between accuracy and the amount of information, computation, or time. This means that system 2 processes and following more classically “rational” decision-making models do not always lead to better decision-making quality. For instance, relying on one heuristic, such as “take-the-best” (Czerlinski et al., 1999), can lead to higher predictive accuracy than that of multiple regression models in environments of moderate to high uncertainty (Hogarth & Karelaia, 2007) and redundancy between mental cues (Dieckmann & Rieskamp, 2007).



**Figure 3** - *Inverted U-shape: “Less-is-more”* (Source: Adapted from Gigerenzer & Gaissmaier, 2011).

To counteract the drawbacks of heuristic decision-making processes and bounded rationality, Stobierski (2019) identifies the importance of identifying and utilizing the correct data to create the foundation of an informed and rational decision. As such, data-driven decision-making requires collecting and analyzing the right data

to ensure that decisions are based on accurate and up-to-date information (Stobierski, 2019). To achieve this, Sarker (2021b) identifies a variety of techniques and methods for collecting and analyzing data, such as data mining, predictive analytics, and artificial intelligence.

### 2.2.2 Organizational and Situational Context

Decision-makers in organizations need to consider what is right to do in the context of the organization, which March (1994) coined as “the logic of appropriateness”. This concerns an individual’s understanding of its role in the organization, what characterizes the situation, the company, and what the organizational member should do in a situation (March, 1994; Simon, 1945). Although this is an ideal decision-making strategy for organizational members, one cannot guarantee that employees will follow this. A well-known example is sub-optimization, in which a part of a company prioritizes its own success to the detriment of the company’s performance (Cambridge Dictionary, n.d.).

Organizational decision-making can be affected by goals and strategies, formal structure, culture, power structures, information technology (IT), personal and situational characteristics (Jacobsen & Thorsvik, 2019; Shepherd & Rudd, 2014), and emotions (Lerner et al., 2015; Simon, 1967). Although we recognize that all of these can significantly influence organizational decision-making, we are mainly concerned with strategy, IT, and situational characteristics, such as uncertainty, complexity, and speed concerns, which are already addressed. IT is naturally discussed in the next sub-chapter, 2.3.

Chandler (1962) defines strategy as “the determination of the basic long-term goals of an enterprise, and the adoption of courses of action and the allocation of resources necessary for carrying out these goals” (p. 13). Central in the strategy research paradigm is the seminal work of Michael E. Porter (1985) on strategic positions and how organizations can achieve competitive advantages by following one of the generic strategies: cost leadership, differentiation, and focus. By choosing one of these strategies, the information search part of a decision-making process will be biased towards said strategy (Cyert & March, 1963). For instance,

if a company has a sustainability-related strategy, it is likely that information search processes will be biased toward sustainability issues and place less emphasis on other aspects of the search process. The more specific and unambiguous the strategies and goals are, the more they will constrain the decision-making of employees and increase the predictability of outcomes (Cyert & March, 1963). Central to strategy is effective resource allocations, which can have important implications for organizations wanting to innovate and implement AI.

### 2.3 Artificial Intelligence in Organizational Decision-Making

More than 50 years ago, Herbert A. Simon (1973), along with other organizational theorists, pointed to the potential of IT to revolutionize how decisions are made in organizations: “He [the human] is more and more an observer... and repairman for a nearly autonomous process that can carry on for significant intervals of time without direct human intervention” (p. 269). Also, through computers’ enormous potential for data collection and processing, it was claimed that organizations could move to a new level of rationality (Simon, 1973). Theorists were very optimistic at the time as the progress in computing technology was rapid, with the likes of Marvin Minsky predicting in 1967 that AI will be “substantially solved” within a generation. Despite the positive outlook, this would take much longer than anticipated (Brynjolfsson & McAfee, 2017).

With the recent developments in AI-based technologies, we are now seeing a revolution in how organizational decisions are made. By partly or fully automating and augmenting decision-making processes, AI can, if developed and implemented correctly, help organizations improve efficiency, reduce costs, increase accuracy, create new information, pursue new opportunities, and make better decisions (Mollick, 2022; Borges et al., 2021; Shrestha et al., 2019; Colson, 2019; Brynjolfsson & McAfee, 2017). In addition, with decision quality approaching human-like intelligence in specific problem domains, AI has become an increasingly popular tool for organizations across a range of industries. The openness and maturity towards intelligent systems, however, varies greatly within organizations. Kolbjørnsrud et al. (2017) found that 46% of top-level managers



would trust the advice of intelligent systems when making decisions, whereas only 14% of first-line managers would. The authors found that the level of trust toward intelligent systems is closely connected with the inherent understanding and familiarity of the system itself. In their study, the authors found that 61% of managers expressed the need to understand how the system works in order to establish trust, and 57% highlighted that a proven track record would allow them to trust in the system (Kolbjørnsrud et al., 2017).

In addition to this, there are significant drawbacks and risks that have been highlighted in previous literature that need to be addressed, in order to achieve a sustainable and responsible use of AI (Shrestha et al., 2019; Castelvechi, 2016; Emanuel & Wachter, 2019; Mathew et al., 2021; Wilks, 2019). According to Jacobsen and Thorsvik (2019), there will always be a need for people who can make discretionary decisions, and that no data system can guarantee good decisions. Evident by NLP models possibility of ‘hallucinating’ and generating factually wrong data (Ji et al., 2023). As such, a data system may be able to provide information, but it does not necessarily have to be good and valid information (Jacobsen & Thorsvik, 2019).

### 2.3.1 Complementarity Between Humans and Machines

Humans excel at making decisions that benefit or require tacit knowledge, or in other words, knowledge that cannot easily be articulated and stored (Cambridge Dictionary, 2023a; Metcalf et al., 2019). Explicit knowledge, on the other hand, is knowledge that can easily be codified and stored (Cambridge Dictionary, 2023b). When conducting a comprehensive analysis, a significant amount of data has to be reviewed in order to find a pattern. As such, humans are likely to experience memory limitations when processing a considerable amount of explicit knowledge (Tversky & Kahneman, 1974). Computers, on the other hand, excel at processing explicit knowledge (Metcalf et al., 2019). Thus, introducing the notion that in combination, when machines and people are connected in the decision-making process, greater intelligence, and higher quality decisions are achievable (Cannavacciuolo et al., 2015; Davenport & Ronanki, 2018; Engelbart, 1962; Harlow, 2018; Schoemaker & Tetlock, 2017; Shrestha et al., 2019).

Shrestha et al. (2019) examined how human decision-making in organizations is affected by AI-based decision-making algorithms and compared them across five decision-making conditions: 1. Specificity of the decision search space, 2. Interpretability of the decision-making process and outcome, 3. Size of the alternative set, 4. Decision-making speed, and 5. Replicability of outcomes. First, a human can exercise tacit judgment and intuition to address ill-structured decision objectives. In contrast, today's AI algorithms are dependent on a well-structured decision space, representing narrow AI. Second, a human can easily explain and justify their decision-making process, although this might not be "accurate, truthful, or comprehensive" and vulnerable to biased retrospective sense-making (p. 4). AI algorithms, conversely, utilize complex optimization techniques to identify patterns in data, but lack explainability, interpretability, and thus, the ability to detect biases, avoid them and generate trust. This is known as "black box" models, in which the process of converting inputs to outputs lacks transparency. Third, AI algorithms can be uniformly applied across millions of alternatives, which is not physically possible for a human to achieve as we would likely be overwhelmed (Iyengar & Lepper, 2000), and have a higher chance of making the wrong choice (Inbar et al., 2011) and leading to a state of paralysis (Langley, 1995). Fourth, as system 2 reasoning and speed have been discussed previously, we will exclude this point from this discussion. However, it is noteworthy that Shrestha et al. (2019) mirror Tversky & Kahneman (1974) and Kahneman's (2011) view on how fast and heuristic decision-making represents fundamental errors, making the argument that when humans are fast, they are often irrational. Lastly, AI algorithms will produce consistent outputs based on consistent inputs, whereas humans lack this replicability because of aforementioned cognitive reasons, such as fatigue.

### 2.3.2 Decision-Making Structures

IT's greatest impact on organizations is the complete or partial automation of some decisions, where several researchers have proposed different frameworks of automation (Colson, 2019; Shrestha et al., 2019; Jacobsen & Thorsvik, 2019). Common to these models is the objective of leveraging the respective strengths of

AI algorithms and human decision-makers to hopefully suppress their respective weaknesses. Based on this, we end up with the following framework (Colson, 2019; Shrestha et al., 2019; Jacobsen & Thorsvik, 2019):

<b>Degree of automation</b>	<b>Decision-making process</b>
Full	Information → AI → Decision(s)
AI to Human	Information → AI → Possible Action(s) → Human Judgment → Decision(s)
Human to AI	Information → Human judgment → Possible Action(s) → AI → Decision(s)
Aggregated AI and Human (Augmentation)	Information → AI ∨ Aggregation rule → Decision(s) Information → Human ↗
Not at all (Assisted by Machine or Not)	Information → Summarized Data → Human Judgment → Decision(s)

**Table 2 - Degree of Automation in Decision-Making Processes.**

In the case of full automation, data is inserted into an AI-algorithm that produces an output, the decision (Jacobsen & Thorsvik, 2019; Colson, 2019). However, it is important to point out that this is limited to routine decisions that rely on structured data, in which we are better off delegating decision-making to AI (Colson, 2019). This means that the decision can be standardized to a set of questions that can give clear and unambiguous answers (Sheridan, 1992). The AI would have no trouble going through millions of groupings in a replicable manner and is also comfortable with nonlinear relationships (Colson, 2019). In addition, this is a favorable structure if the accuracy and speed of the prediction are more important than interpretability, as there typically is a trade-off between these (Shrestha et al., 2019). This could for instance be the case in high-speed environments, such as high-frequency trading where speed is essential (Shrestha

et al., 2019). Also, it is worth noting that this structure comes with several limitations. This can include that if the patterns in the data were to change over time, the algorithm would be less accurate, in addition to significant ethical issues (Shrestha et al., 2019), which will be discussed later in this chapter.

As there are many decisions that go beyond the scope of structured data, humans can complement AI by leveraging other types of information unavailable to computers (Colson, 2019). Such information can be related to strategies, values, and market dynamics that cannot be transferred through digital communication (Colson, 2019). These hybrid models can be performed through human-to-AI, AI-to-human, or aggregated AI and human decision-making processes, with each model suited to different contexts (Colson 2019; Shrestha, 2019). In these ways, one can combine humans' specialized input with that of the "objectively rational" data-processing computers, to achieve better decisions than either of them could separately (Colson, 2019). The contingencies for AI automation as previously mentioned apply to AI in the hybrid models as well, with the key differences being that humans increase the interpretability of a decision, but at the expense of replicability and speed (Shrestha et al., 2019).

In the case of not utilizing AI, computers are still used, but to a limited extent or not at all (Jacobsen & Thorsvik, 2019). In the case of using computers, or a "data-driven" workflow, large amounts of data are summarized through databases, spreadsheets, et cetera (Colson, 2019). However, as with the other decision-making structures, this model has its limits. This includes the obscuration of "insights, relationships, and patterns contained in the original [big] data set", which makes this process somewhat invaluable, and perhaps misleading, to decision-making processes (Colson, 2019). Also, with humans as the central processor, the summaries are prone to cognitive biases.

### 2.3.3 Algorithmic Bias, Transparency, and Data Issues

Decisions and choices that traditionally have been left to humans are increasingly being delegated to algorithms (Mittelstadt et al., 2016). As the use of machine learning models and AI tools is increasing in popularity, so are the concerns

related to algorithmic bias and transparency issues (Shrestha et al., 2019). With algorithms being a part of business transactions, social processes, and decision-making, to name a few, our perception and understanding of the environment we operate are changing (Mittelstadt et al., 2016).

Despite significant advances in deep learning, which NLP is based on, over the past decade, the technologies continue to face certain limitations that constrain their full potential. One of the most salient and pressing of these limitations pertains to the conspicuous lack of interpretability that characterizes deep learning models (Castelvecchi, 2016). This limitation arises primarily due to the massive size and complexity of contemporary deep learning models, which consist of trillions of parameters. Consequently, comprehending how these models arrive at their predictions has become increasingly challenging, if not nearly impossible, for humans. As a result, deep learning models are often referred to as black boxes. This lack of transparency can make it difficult to trust the results of these models (Castelvecchi, 2016; Shrestha et al., 2019). Although there exist models that “open the black box” by making the non-linear and complex decision processes humans, these models currently offer limited performance compared to their opaque ancestors (Eschenbach & Warren, 2021).

Castelvecchi (2016) argues that machine learning will never be explainable because a machine able to understand the real world will be complex, simply because the real world itself is complex and there are things that cannot be verbalized. Stéphane Mallat (referenced in Castelvecchi, 2016) explains the rationale behind this as: “When you ask a medical doctor why he diagnosed this or this, he's going to give you some reasons.... But how come it takes 20 years to become a good doctor? Because the information is just not in books.... You use your brain all the time; you trust your brain all the time; and you have no idea how your brain works.” (p. 9).

Another limitation is the bias in these models. As they are trained on large datasets consisting of online textual data, which may contain biases toward race and gender that can be inadvertently learned and replicated by the models (Bolukbasi et al., 2016). In other words, a model trained on a biased dataset may

associate certain words with specific demographic or cultural stereotypes. An example is a popular online translation system that utilizes statistical machine translation, which has been documented to construct gender-stereotyped translations from gender-neutral languages (Shrestha et al., 2019). Shrestha et al. (2019) suggest that AI-based decision-making may not only perpetuate cultural stereotypes and discrimination, but also amplify them.

In addition to this, predictive analysis has several limitations as well (Kumar & Garg, 2018). Overfitting is one of the most common problems in machine learning. This occurs when the model is too complex and fits the training data too closely. This can lead to poor performance when applied to new, unseen data (Trivedi et al., 2021). Another limitation is the potential for bias, which occurs when the resulting model is skewed and systematically off-target in a particular direction. This can occur when the data set used to train the model does not represent the population being studied, usually due to human intervention, chance, or poor data collection (Kelleher et al., 2020; Lones, 2021). The quality of the data itself is another significant limitation, as the accuracy of the predictions is limited by the quality of the data used. If the data is noisy, incomplete, or inaccurate, the resulting predictions will also be of poor quality; a process that is generally known as a “garbage in garbage out”-model (Lones, 2021).

#### 2.3.4 Organizational Trust in AI

Trust can be defined as “an optimistic expectation on the part of an individual about the outcome of an event or the behavior of a person.” (Hosmer, 1995, p. 390). According to Galford and Drapeau (2003), trust in organizations consists of three different types: 1. Strategic trust, which is the trust that the top-management team makes sound strategic choices to make sure that the set course is aligned with the organization’s visions and long-term goals, and the ability to allocate resources intelligently, 2. Personal trust, which is the trust between individuals in the organization, and 3. Organizational trust, which is the trust in the organization itself- that routines and processes are well designed, consistent, and fair. While we recognize that these are three distinct types of trust, they are intricately connected. As such, to build and maintain trust within an organization, it is essential to

manage these three aspects of trust properly (Galford & Drapeau, 2003). Jacovi et al. (2021) further states that there are two types of trust towards AI tools, namely intrinsic and extrinsic trust. The former can be gained through the explanation of a model, or in other words, transparency and explainability of the outcome. The latter, however, can be obtained through the evaluation of data, or in other words, consistently accurate and therefore trustworthy outcomes.

There are two different methods of obtaining intrinsic trust: 1) the user successfully comprehends the true reasoning process of the model, and (2) the reasoning process of the model matches the user's priors of agreeable reasoning (Jacovi et al., 2021). To obtain extrinsic trust, on the other hand, there are two methods: 1) By proxy, and 2) By observation. In other words, a user can obtain trust towards an AI tool by being influenced by an expert human opinion, or by observing that a model produces consistently accurate outcome over time in different contexts.

### 2.3.5 Successfully Implementing AI

According to a study conducted by Deloitte, 94% of business leaders agree that AI is going to be critical for success over the next five years – yet the actual implementation is lagging behind (Deloitte, 2022). While the artificial intelligence readiness index rankings vary greatly between countries, The United States of America is the highest-ranked country with an index score of 85,72. Norway, on the other hand, is on 12<sup>th</sup> with a score of 73,09 (Deloitte, 2022). In other words, the average company seems to be ready to utilize AI solutions. According to the latest report about the state of AI from McKinsey, the share of respondents who say their organizations have adopted AI in at least one business unit or function has increased from 20% in 2017 to 50 % in 2022, where the majority of use cases stems from service operations, i.e., optimization, and the creation of new AI-based products (McKinsey & Company, 2022).

Despite offering several advantages, many organizations are hesitant to develop and implement AI, which could explain why only 50% of organizations in OECD have adopted it (Lane et al., 2023). According to Deloitte's (2022) study, there are

three overall challenge areas in scaling AI initiatives. In the first area, barriers to starting projects, 50% of the participants highlighted the difficulty of managing AI-related risks, and 44% highlighted challenges connected to obtaining training data to develop a model. In the next problem area, implementation, 50% of the respondents highlight insufficient executive commitment, while 42% highlight alignment difficulties between AI developers and business needs/problems. In the last problem area, continued scaling, 46% of the respondents highlighted difficulties connected to integrating AI into an organization's daily operation as the biggest challenge, and 44% highlighted that the developed AI solutions were too complex and difficult for end users to adopt (Deloitte, 2022).

Similarly, Leonard-Barton & Kraus (1985) emphasize the importance of starting with user needs and preferences. This approach encourages early involvement by users in the design phase to enhance the fit between their needs and the new technologies, boosting user satisfaction and increasing collaborative communication. The authors also see this as a vital part of implementation success, demanding sustained level of investments, not only throughout the initial phases of development, but also during the often less prioritized implementation phase. They also discuss the role of new technology hype, in which expectations can far exceed performance. This stresses the importance of not overselling the technologies, but at the same time not underselling it either. The AI label can be particularly misleading as it has grown synonymous with human-level capabilities and artificial general intelligence, which contributes to artificially high failure rates for ML initiatives (Siegel, 2023). In addition, it is the more simple and practical use cases of AI “that deliver the greatest impact on existing business operations” (Siegel, 2023). This results in recommendations about resisting the temptations of the AI hype wave.

Brynjolfsson and McAfee (2014) found that the successful implementation of AI in organizations is largely dependent on strategic alignment, sufficient resources, and if the organizational culture is ready to implement and adopt AI. Furthermore, the authors argued that an organization needs to ensure that the deployment of AI aligns with its overall strategic objectives and business model.



In their three-step model, Kolbjørnsrud et al. (2017) found that the first step towards implementing AI entails exploring AI internally in the organization so that the employees can develop a better understanding of what AI is and how it works before implementing AI tools in the organization. As such, developing a level of trust towards the AI, in which both human and digital actors need to be able to account on each other for carry out their specific tasks consistently (Kolbjørnsrud, 2023).

Bughin et al. (2017) suggested that fostering an AI-friendly culture that promotes continuous learning and adaptation is vital for successful AI adoption. Their findings underscored the need for employees at all levels to understand and engage with AI, necessitating ongoing training and development initiatives. This aligns with Kaplan and Haenlein (2019) and Kolbjørnsrud's (2023) emphasis on the importance of managing human-machine interactions effectively to reduce resistance to AI and foster acceptance. The former authors further recommended that leaders communicate transparently about the benefits and limitations of AI, addressing employees' concerns, which can enhance their general understanding and acceptance of AI.

Furthermore, several scholars have highlighted the importance of investing significant resources in data, technology infrastructure, and maybe the most important one, human skills to harness the benefits of AI (Brynjolfsson & McAfee, 2014; Davenport & Ronanki, 2018). Kolbjørnsrud et al. (2017) further highlight the need for managers to emphasize the need to recruit employees with soft skills such as collaboration, creativity, and good judgment, which is just as, or even more important than recruiting employees with technical skills. The reason being that such qualities complement and augment the more analytical qualities offered by intelligent systems (Kolbjørnsrud et al., 2016).

### 2.3.6 The Existing State of Literature

To sum up the literature review, there is no perfect way to utilize AI in organizational decision-making processes. Issues related to transparency and the black box challenge and how it affects decision-making processes remain some of

the most prominent reasons. Furthermore, the trade-off between transparency and performance introduces unanswered questions, such as under what circumstances do organizations require transparency of reasoning and how can this trade-off be managed? (Dwivedi et al., 2021). Although the existing literature has focused on one side of this issue, namely on the technical advancements and the development of explainable methods, less attention has been devoted to the other side and the impact these methods have on end-user's trust, understanding, and their ultimate decision-making (Burrell, 2016). To what extent can complex machine learning models truly become acceptably interpretable? How does the interpretation provided by these models align with human cognitive processes and intuitive understanding? Furthermore, to deal with the concern connected to amplifying existing biases in a data set (Shrestha, 2019), how can developers ensure that the quality of the data is sufficient to support the required analysis?

While several ethical frameworks for AI have been proposed to manage the ethical concerns (Jobin et al., 2019), current literature often treats these ethical considerations in a somewhat abstract manner (Hagendorff, 2020; Mittelstadt, 2019; Morley et al., 2021). Few studies have addressed their practical application and effectiveness in real-world decision-making settings, which suggests that there is a clear need for more empirically grounded research that explores these issues within real-world organizational contexts. That is, how is a model managed if it makes a mistake? Who is the responsible and accountable party in cases of incorrect or harmful decisions; the human who executed them, or the AI model that suggested them? (Cath, 2018). This further introduces how significant the role of trust is and how it relates to AI and decision-making. And extending this, how can trust and accountability be achieved as AI becomes increasingly important in decision-making? (Shrestha et al., 2019). How do the different stakeholders deal with ethics and trust? (Dwivedi et al., 2021).

Davenport and Kalakota (2019) have examined AI's role in streamlining decision-making processes, predominantly focusing on the operational aspects. Their research does not fully explore how AI augments the decision-making process itself, beyond simplifying data analysis and providing more efficient workflows.

Although there are studies on AI's capability to provide predictive analytics and how it can influence decision-making (Agrawal et al., 2018), they fail to address how AI could improve the decision-making process by providing fresh insights or perspectives that humans may overlook due to cognitive biases. For instance, AI's capability to discover non-obvious patterns in large datasets can help organizations make decisions in complex environments (Bughin et al., 2017). The issue, however, is that the current literature does not sufficiently investigate how this capability translates into real-world decision-making scenarios.

Furthermore, the relationship between how AI and culture influences decision-making seems to be unexplored (Dwivedi et al., 2021). This raises unanswered questions, such as how organizations should structure their business and technology architectures to support the integration of new technology. In addition, further research is needed on how performance should be measured in AI-augmented decision-making, how different AI decision-making structures impact organizational performance, and how the decision-making context prescribes what structures are suitable (Shrestha et al., 2019).

To increase our understanding of how AI enhances organizational decision-making processes, we aim to address some of the existing gaps and unanswered questions. There is a lack of empirical studies on how to create business value with the adoption of AI technologies (Borges et al., 2021, Brynjolfsson & McAfee, 2017; Dwivedi et al., 2021) and there is limited research on AI employing cross-case methodology (eg. Ferrer et al., 2021; Trocin et al., 2021). Responding to these issues, our study is grounded in the belief that understanding the full spectrum and dynamics of AI's impact on decision-making requires not just a technical understanding, but also a deep exploration of its integration and implications across industries and technologies within an organizational context. This is in line with Engelbart (1962), that the augmenting system of humans and machines “can best be improved by considering the whole as a set of interacting components rather than by considering the components in isolation.” (p. 2). As such, we aim to discover how organizations can effectively integrate AI into their organization and address the intricacies between benefits and challenges.

## 3.0 Methodology

After reviewing the literature, we conducted a multiple case study to get empirical evidence on the most central factors for how AI augments organizational decision-making. In this section, we will present and explain the methodological choices we have made to effectively address the research question.

### 3.1 Research Design

#### 3.1.1 Qualitative Research Design

To answer our research question, we have opted for a qualitative methodology approach, which is particularly suitable for investigating phenomena where there is limited knowledge and scant research (Yin, 2009). Although the introduction of new technology in organizations is not a new phenomenon, the adaptation of artificial intelligence in organizations, however, is a relatively new notion for many companies. As such, the existing literature on how AI augments organizational decision-making processes is relatively scarce. As a result, we deemed qualitative methodology to be the most suitable research method for this case.

Furthermore, we have chosen case design as our method of study. Denscombe (2014) defines case design as a process involving the formulation of a research question, selection of a case, choice of informants, data collection, and criteria for data analysis and interpretation. They further state that case studies provide an in-depth examination of events or organizations to investigate phenomena that might otherwise remain undiscovered. As such, this approach was a natural choice based on our research question and the topics to be covered. This allowed us to delve deeply into each organization and observe how AI impacts the organization as a whole, in addition to AI's impact on a single-employer basis.

We employed an exploratory design as our starting point, which is suitable when there is limited prior knowledge about the subject of study (Creswell, 2014), as was our case. We were cognizant of the changes that must have occurred in most

industries but were unaware of how management had evolved in the consulting industry over the past year, as well as the exact changes that had transpired. For exploratory studies, Yin (2009) suggests having a clear and defined purpose for the study, rather than a pure theoretical assumption. This approach allowed us, as researchers, to maintain an open perspective on the study and gain better insight into how AI augments the decision-making process in an organization. Despite the open and exploratory procedure, we used the theory as a basis for analyzing and discussing our results, in addition to initially identifying areas and topics with limited academic coverage.

We opted for a multiple-case design with 11 analysis units, or cases, a so-called cross-case analysis. This design facilitates comparison to the extent that it is possible to compare the two departments. We will revisit this towards the end of the paper in the subchapter on potential limitation of the study. An analysis unit can be defined as "a social unit or the element in society that the study is based on." (Grønmo, 2004, p.79). The primary advantage of cross-case analyses is that it enhances the relevance and transferability to other settings, potentially contributing to a higher degree of generalizability (Yin, 2009). It is worth noting, however, that generalizability is not a central criterion in qualitative research, where the focus is often on selecting a limited number of informants who possess extensive and relevant information about the phenomenon under investigation (Johannessen et al., 2020).

The interview guide we developed for this study was predominantly exploratory and open-ended. As such, we began the interviews by asking open and descriptive questions and then following up with deeper questions exploring opinions and personal experiences. Our objective was to adopt an unbiased perspective concerning the impact of AI on decision-making processes, and the perceptions of the interviewees. This guide was also grounded in the theoretical component of our research, as one of the aims was to observe how theoretical expectations compare with empirical analysis and outcomes. For accuracy and quality insurance, we discussed the layout and content of the guide with our supervisor before conducting the interviews. For a comprehensive interview guide, please see

the Appendix. In addition to interviews based on the guide, we asked follow-up questions tailored to each case where necessary.

### 3.1.2 Case Sampling

The recruitment process was primarily conducted through our own network, utilizing the snowball method, which enabled us to identify and select supplementary interviewees to get a better understanding of each case (Bryman & Bell, 2015). Our aim was to interview 3-4 employees in each case that are either managing the implementation of AI solutions, developing operational AI solutions, or are being impacted by it, in a broad spectrum of industries. In order to maximize variation sampling within a specific topic, we chose to include cases from several industries in our study. This allowed us to observe and enhance our understanding of the phenomenon at hand more accurately (Suri, 2011).

Regarding the type of decision, we were interested in recurring decisions where patterns could potentially emerge. Before embarking on comprehensive interviews, we were receptive to both strategic and operational uses of AI in the decision-making process, and cases with different degrees of success. The endeavor to identify and come in contact with potential candidates proved to be more challenging than anticipated. As a result, we opted to conduct up to several in-depth interviews with each interviewee in cases where we were not able to achieve our initial goal. We do, however, believe that this has not affected the accuracy or generalizability of our study significantly. Nonetheless, it is worth emphasizing that in qualitative studies, the goal is not necessarily to achieve generalizability (Johannessen et al., 2020). Instead, we view this as a strength in terms of obtaining diverse perspectives on the research question at hand.

In total, we interviewed 23 people distributed across 11 cases. Of these 23 participants, 10 were data scientist leads or data consultants, with varying degrees of experience. The remaining 13 were either business managers, business consultants, or employees on the operational level in the firms. By combining both the technical-oriented and business-oriented aspects of AI in decision-making processes, we were able to maximize variation sampling within a specific

topic, which we believe enabled us to accurately observe how AI augments the organization as a whole (Suri, 2011).

Table 3 summarizes the number of participants and the number of interviews from the different cases included in our study. The symbol in the top right corner symbolizes if the case was considered to be successful or not.

<b>Cases With Predictive Analytics</b>			
<b>Case 1: Fleet Management</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 1	2	Head of Data and AI	Transportation
<b>Case 2: Deviation Prediction</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 2	2	Chief Commerical Officer	Water
<b>Case 3: Transport Modelling</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 3	1	Chief Data Scientist	
Interviewee 4	1	Head of Digital Mobility	Transportation
Interviewee 5	1	Assosicate Data Scientist	
<b>Case 4: Water Level Management</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 6	1	Head of Digital and Innovation	Water
<b>Case 5: Maintenance Offshore Wind</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 7	1	Management Consultant	Offshore Wind
<b>Cases With Natural Language Processing</b>			
<b>Case 6: CV &amp; Job Matching</b>			
			✗
Interviewee	Number of interviews	Position	Industry
Interviewee 8	3	Data Scientist	
Interviewee 9	1	Office Manager	Recruitment
<b>Case 7: Policy Comments Review</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 10	1	Data Science Consultant	IT Consulting
<b>Case 8: Bid Robot</b>			
			✓
Interviewee	Number of interviews	Position	Industry
Interviewee 11	1	Director of Digital Business Developme	IT Consulting

<b>Cases With Classification Analytics</b>			
<b>Case 9: Bacteria Detection</b> ✓			
Interviewee	Number of interviews	Position	Industry
Interviewee 12	1	Head of Digital and Innovation	Water
Interviewee 13	1	Management Consultant	
Interviewee 14	1	Senior Consultant	
<b>Case 10: International Customer Transactions</b> ✓			
Interviewee	Number of interviews	Position	Industry
Interviewee 15	1	Chief Data Officer	Finance
Interviewee 16	1	Data Scientist - Robotic and AI	
<b>Case 11: House Insurance</b> ✓			
Interviewee	Number of interviews	Position	Industry
Interviewee 17	2	Team Leader - AI and ML	Insurance
Interviewee 18	1	Customer Service Agent	
Interviewee 19	1	Project Manager - IT	
<b>Introductory Landscape Interviews</b>			
<b>Other</b>			
Interviewee	Number of interviews	Position	Industry
Interviewee 20	1	Director of Business Development	Insurance
Interviewee 21	1	Chief Technology Officer	Technology
Interviewee 22	1	Head of Automation	Water
Interviewee 23	1	Middle Manager	Water

*Table 3 - Overview of Included Cases*

### 3.1.3 Case Descriptions

As previously stated, we have, in this study, opted for a multiple case study design, as this provides a more comprehensive and balanced perspective on the interplay between AI and organizational decision-making processes. Out of our 11 cases, 5 of these are based on predictive analytics, 3 are based on classification analytics, and the remaining 3 are based on natural language processing. This section of the thesis will provide a short summary of each case, and briefly explain why we have included them in our study. In addition to this, several of the interviewees have referred to additional cases during the interviews. We have included some statements from these in our analysis, but do not provide a description of these below. We believe the inclusion of these has strengthened the breadth, accuracy, and validity of our data foundation, by maximizing variation sampling within a specific topic (Suri, 2011). The role as ‘operator’ as referred to throughout the study is the individual at the end of the process that is using and acting on the output from an AI tool.



### 3.1.3.1 Cases with Predictive Analytics

#### **Fleet Management**

This case was about a fleet management service where a machine learning model was trained to predict where new orders were likely to originate based on historical data. By dynamically placing their minibuses closer to the next customer, the company was able to shorten wait times and reduce negative environmental impact by limiting unnecessary driving and idling time. By including this case in our study, we aimed to illustrate how an organization can utilize machine learning models in decision-making processes by providing data-driven insights for resource allocation.

#### **Deviation Prediction**

This case was about deviation detection, in which a predictive machine learning model was trained to identify the early signs of potential issues in water treatment plants or their connected piping infrastructure. The model used measurement values from various plants to predict whether the development in a process is positive or negative based on historical data. It was also able to continuously update the normal limits and adjust the acceptable limits. By identifying deviations before they become critical problems, the model enabled an operator to take preventive measures and avoid system failures before they happened, ensuring system reliability, and minimizing downtime. Thus, it enabled the operator to make more informed decisions when choosing how to maintain and respond to internal alarms. By including this case in our study, we aimed to illustrate how the operator is affected by including predictive machine learning in the decision-making process.

#### **Transport Modelling**

This case was about a consultancy firm that developed an AI-driven activity-based travel demand modeling software to provide comprehensive analysis and insights into the demographics and socioeconomic impacts of various transportation measures in a medium-sized city. It differed from traditional transport models by incorporating synthetic populations representing real individuals and their behavioral choices in their daily life. By leveraging machine learning models to

analyze both real and simulated transportation data, the software was able to more accurately illustrate how potential changes in public transport, roads, and pathways would impact the mobility flow in the city. By including this case in our study, we aimed to illustrate how an organization can utilize machine learning models in decision-making processes by providing data-driven insights when planning resource allocation and identifying focus areas to improve efficiency.

### **Water Level Management**

This case was about a consultancy firm that utilized a machine learning model to enhance the decision-making in water management by accurately predicting water demand, optimizing supply, and balancing cost efficiency and environmental sustainability. By leveraging machine learning models to analyze historical water consumption data, weather information, and calendar events, the company was able to predict the water demand of the population for the next day and use this information to make informed decisions about the amount of water to pump into water towers at night when electricity cost is the lowest, and whether to purchase additional water from external sources. By including this case in our study, we aimed to illustrate how the system operators are affected by including predictive machine learning in the decision-making process.

### **Maintenance of Offshore Windmills**

This case was about a consultancy firm within the offshore wind sector that developed a machine learning model to predict the required maintenance needs of a large network of individual windmills, based on a small dataset consisting of sensor data from just a few windmills. Traditionally, each windmill has been required to be manually inspected for damage or wear. By leveraging machine learning models to analyze sensor data from a few windmills, the company is able to predict future downtime and damages to the wind farm as a whole, and as such, offer a cost-effective alternative to monitor windmills with an accuracy of 90% without equipping the entire network with expensive sensors. By including this case in our study, we aimed to illustrate how the maintenance operator is affected by including predictive machine learning in the decision-making process.

### 3.1.3.2 Cases with Natural Language Processing

#### **CV & Job Advertisement Matching**

This case was about a governmental department that utilized NLP models to assist hiring managers when choosing which candidates to interview for an open job position. Traditionally, this has been a manual and time-consuming process done by reading hundreds of applications and CVs per position. By utilizing NLP models trained on newspaper articles, the department was able to rate the relevancy of a CV to a certain position by comparing text similarity. With an initial accuracy rate of 75%, the hiring manager was able to better and significantly quicker match relevant CVs to a certain position. By including this case in our study, we aimed to illustrate how the hiring manager is affected by including NLP models in the decision-making process.

#### **Policy Comments Review**

This case was about a consultancy firm that utilized NLP models to analyze and summarize large amounts of qualitative data. When evaluating feedback about future a policy change or a product, one has to consider thousands of pages with comments and texts in order to accurately represent the feedback at hand. This has been, however, time-consuming, and nearly impossible for humans due to the sheer amount of data. By utilizing NLP models the company was able to analyze and summarize large amounts of feedback, and as a result, provide a foundation for well-informed decisions. By including this case in our study, we aimed to illustrate how the firm managed to enhance its decision-making accuracy by including NLP models in the decision-making process.

#### **Bid Robot**

This case was about a consultancy firm that utilized NLP models to assist project managers when developing tenders and identifying risks in these, for potential customers. Before a consultancy firm can take on a new project for a client, developing a tender is the first step. This is, however, a time-consuming and critical process in terms of risks and responsibilities later in the project. By utilizing NLP models trained on legal language and risk classifications, the company was able to aid project managers by automatically identifying risks in a

tender, and as such, ensuring that the tender is meeting the requirements of the customer at the same time as protecting the consultancy firm from potential legal claims and lawsuits in the future. By including this case in our study, we aimed to illustrate how the project manager is affected by including NLP models in the decision-making process of whether to send the tender to a potential customer or not.

### 3.1.3.3 Cases with Classification Analytics

#### **Bacteria Detection in Water**

##### **Clearwater:**

This case was about a start-up that used a machine learning model to analyze and monitor algae blooms in lakes by capturing a photo on a portable field microscope. Traditionally, this process has been done by volunteering citizens that collected water samples and sent them for thorough lab analysis by trained scientists. This information has then been used to decide on whether to implement actions, such as closing down a beach. By leveraging machine learning models to analyze pictures instead of actual water samples, the automation of the analysis section provides a significant reduction in lag time, and made it cheaper to process a sample, as it does not require trained scientists to operate. As such, making it accessible to citizen volunteers or low-paid field workers on a larger scale. By including this case in our study, we aimed to illustrate how the process is affected by including machine learning models in the decision-making process.

##### **Wastewater:**

Using similar technology as in the clear water case, this case was about a water treatment plant that utilized machine learning models to analyze and monitor bacterial growth in wastewater treatment tanks. By accurately analyzing and assessing the presence and severity of issues, the wastewater plant was able to significantly reduce the risks of overflows and contamination. By including this case in our study, we aimed to illustrate how the system operator is affected by including machine learning models in the decision-making process of how to manage wastewater treatment.

### **International Customer Transactions**

This case was about a bank that used a machine learning model and NLP to categorize and automate the approval of cross-border transactions with other foreign banks. Traditionally, the approval of a transaction has been a manual and time-consuming process. By leveraging machine learning models to understand and categorize free text, the automation of end-to-end processes became possible, allowing for quicker response times and increased efficiency, in addition to identifying potential money laundering cases. By including this case in our study, we aimed to illustrate how case managers are affected by including machine learning models in the decision-making process in complex and fragmented environments, such as with cross-country bank transactions.

### **House Insurance**

This case was about an insurance company that used machine learning models to automate risk assessment by incorporating both standard house data and customer-specific attributes. Traditionally, the process to approve an insurance plan for an old house has been a time-consuming and manual process done by an insurance agent. By utilizing machine learning models to automate the approval process, the insurance company was able to faster and more cost-effective process the approval of new customers. Thus, it enabled the insurance agent to better utilize their time on more difficult cases where the model had flagged a house as not approved. By including this case in our study, we aimed to illustrate how the insurance agent is affected by including predictive machine learning in the decision-making process.

## **3.2 Data Collection**

To gather data, in-depth interviews were conducted, utilizing an open-ended, semi-structured interview guide. As we interviewed different types of employees, both managers, developers, and operators, we used follow-up questions to further investigate their knowledge and experience in their respective areas of expertise. For the interviews with managers, our plan was to prioritize queries regarding their leadership roles in the implementation of AI and the impact on strategic decision-making. For the interviews with developers, we prioritized delving into

the development of the tool, prominent issues, and technical details. For the interviews with operators, or the user of the solution, the goal was to uncover their level of comfort with the solution, its effectiveness in aiding their decision-making where relevant, and its potential for enhancing their valuation of the product or service. Furthermore, we received additional data from some interviewees in the form of 2 PowerPoints and a Ph.D. article.

Given that our participants were dispersed globally, the decision to conduct remote interviews via Microsoft Teams was a logical one. Johannessen et al. (2020) found that online interviews yielded nearly as much invaluable data as their face-to-face counterparts. As a result of the global approach, we conducted interviews in Norwegian, Swedish, and English. Furthermore, we made audio recordings of the interviews to streamline the data collection process. Prior to the interviews, the participants agreed to the interview being recorded and transcribed. Straight after each interview, the interview was transcribed using AI-based transcription services. This output was further quality ensured by comparing the written output to the recording of the interview. We did not alter or modify the transcriptions containing what each interviewee said in any way other than correcting inconsequential grammatical inaccuracies. Overall, the transcription process was relatively unproblematic and resulted in 210 single-spaced pages in Microsoft Word.

We analyzed and codified each transcription in its original language in order to reduce the risk of losing or altering text when translating to English. The main findings and quotes were translated into English so that we could use them when writing our thesis. When translating, we made sure that the translation would be as accurate to the informant's original point of view as possible.

### 3.3 Analytical Process

We adopted a cross-sectional study to examine how AI augments the decision-making processes in our cases. This was done to comprehensively analyze the data we obtained through interviews and other collected materials. This approach entails segmenting the material based on emergent themes and patterns inherent in

the data, without strict adherence to a predetermined theoretical framework. This approach allows for a more flexible and open analysis, where the researcher can let the data speak for itself and uncover new insights and perspectives (Johannessen et al., 2020). Nonetheless, this necessitates a meticulous and systematic approach to ensure the constructed categories are both reliable and valid. To accomplish this, we utilized thematic analysis, a method defined by Braun and Clarke (2006) as a method for identifying, analyzing, and interpreting patterns of meaning or 'themes' within qualitative data. We selected this method due to its flexible nature in handling intricate and nuanced data (Braun & Clarke, 2006).

The first step in this analytical process was to familiarize ourselves with the collected data. This stage involved several comprehensive readings of the transcripts by both authors, a process aligning with Creswell (2014) recommendations to fully understand and comprehend the collected data.

The next stage involved the development of initial codes from the transcripts, through a process known as open coding (Charmaz, 2014). This involved identifying and annotating keywords, phrases, or sentences that seemed particularly significant. Afterward, these codes were grouped into broader categories, which were then developed into themes, through a process known as axial coding (Charmaz, 2014). As we utilized a semi-structured interview guide, the transcriptions were already organized into broad topics and themes. This method, however, allowed us to identify smaller categories and themes and ensure that we didn't overlook any relevant information and statements that lay outside the predetermined structure.

To ensure the trustworthiness and reliability of the findings, we incorporated a form of analyst triangulation, in which the other researcher reviewed the codes and themes derived from the data (Creswell & Poth, 2018). This helped to mitigate potential bias and strengthen the validity of the analysis (Patton, 2015). The first tier entailed a review of the coded data extracts, whereas the second tier revolved around the relation of these themes in relation to the entire dataset (Braun & Clarke, 2006). This approach ensured that the themes were not only

internally coherent, consistent, and distinctive, but also accurately representative of the dataset (Braun & Clarke, 2006).

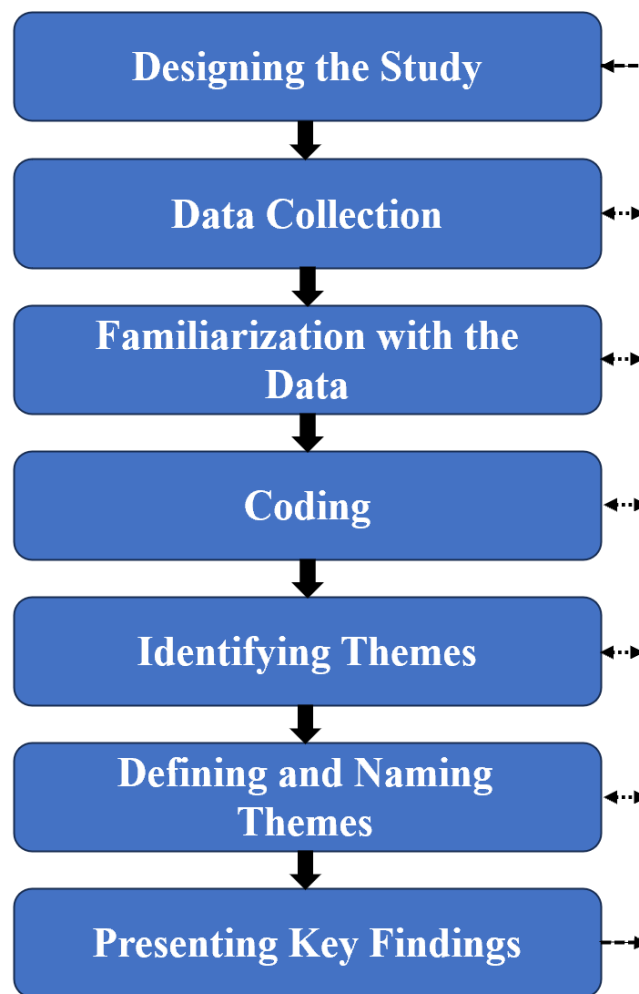
In the next step, we initiated a cross-case analysis, as suggested by Eisenhardt (1989), to identify recurring themes throughout our 11 cases. We implemented a manual color-coding system of the transcripts in Microsoft Word to identify and uncover recurring patterns, thereby facilitating effective comparisons across the cases.

The final stage of the analytical process revolved around synthesizing and interpreting the themes within the context of the research question. This process was guided by Smith et al. (2009) Interpretative Phenomenological Analysis approach, which emphasizes the interpretation of themes in the context of the participant's personal experiences. This involved a comprehensive analysis of how the themes interrelate, what they revealed about the participants' experiences, and how they answered the interview questions.

By maintaining a recursive analytical process, we were able to capture the nuances and depths of the dataset, which reinforces the credibility of our findings (Creswell, 2014). By adhering to these organized yet flexible steps of thematic analysis, we are confident that we have accurately interpreted the meanings, experiences, and perceptions of the participants in alignment with the research objectives.

Table 4 below illustrates our complete analytical process, from designing our study in the beginning, to presenting our finding at the end. As illustrated by the dotted lines on the right side, this was an iterative and partially cyclical process, where we could easily move between, and alter different stages, depending on the need.





*Figure 4 – The Analytical Process.*

### 3.4 Ensuring Quality in the Study

In order to enhance the trustworthiness and reliability of our research, we adhered to rigorous methodological strategies. As such, several measures were adopted to ensure the integrity of our data and findings. According to Bryman & Bell (2015), establishing reliability in qualitative research necessitates the demonstration of consistency and transparency in the research processes. In light of this, we employed member checking, which involved sharing our findings and case understanding with the participants for their verification and feedback. This strategy enhanced the dependability of our study, as recommended by Lincoln and Guba (1985), as it ensured that our interpretation of the participant's responses was accurate. Our initial plan was to interview different people in an organization

to achieve triangulation of data, which allowed us to cross-check data from multiple sources to validate our results (Denzin, 2012). This qualitative method provided a comprehensive perspective and helped to prevent undue bias from any single data source.

In conjunction with dependability, credibility is an integral aspect of qualitative research trustworthiness. Lincoln and Guba (1985) advocate for prolonged engagement with the research context and participants to ensure in-depth data collection and accurate interpretation of the participants' experiences. In this respect, we spent considerable time interviewing participants and interpreting the interview data. This helped us to fully understand the phenomenon under investigation and thus bolstered the credibility of our study. To ensure the authenticity of our research, we endeavored to present the findings in a manner that accurately represented the participants' perspectives and experiences. This entailed accurately transcribing interviews and including direct quotes from participants in the thesis to ensure that their viewpoints were accurately represented (Yin, 2009).

Finally, we paid close attention to the validity of our research. We achieved this by ensuring our research questions, study design, and methods of data collection and analysis were aligned and suitable for our study's purpose (Bryman & Bell, 2015). We also engaged in reflective practices throughout the research process, continually questioning our assumptions and biases to ensure our findings were as objective and accurate as possible (Maxwell, 2013).

### 3.5 Ethical Considerations

When preparing and writing this master thesis, ethical considerations were always a top priority. Our approach was informed by guidance from scholars such as Creswell (2014), who underscored the importance of respect for persons, beneficence, and justice in the pursuit of research.

As our study involved interviewing individuals and recording their responses for further analysis, we secured consent from each participant, acknowledging their agreement to be recorded and transcribed. Throughout the interview process, we

emphasized the voluntary nature of their involvement and their right to withdraw at any stage without reprisal. We were also transparent about the purpose and scope of the research, and any potential implications for the participants were clearly communicated.

Regarding data storage and management, we complied with the guidelines set forth by the National Research Ethics Committee (The National Committee for Research Ethics in the Social Sciences and the Humanities, 2022). As such, collected data was only accessible by the two authors of this thesis, and was stored securely to prevent unauthorized access or data leakage. As within the parameters of consent provided by participants, we are to delete sensitive collected data upon completion of this thesis.

Lastly, during the research reporting process, we were careful to preserve the anonymity of our participants by using pseudonyms and removing identifiable details from the case descriptions and findings. In addition to this, we strived towards maintaining honesty and transparency by acknowledging all contributors and avoiding any form of plagiarism.

We believe that the chosen measures ensured that our research process was ethically sound, respecting the rights and welfare of our participants, while at the same time maintaining the integrity of our scholarly contribution.

## 4.0 Findings and Analysis

In this section we outline the results and analysis derived from the 27 interviews we have conducted across 11 different cases in various sectors. Based on our findings, we have identified 4 major themes, which are: 1. The Perceived Definition of AI, 2. Attitudes and Expectations toward AI, 3. The Benefits of Using AI in Organizational Decision-Making, and 4. The Challenges of Using AI in Organizational Decision-Making. As we also asked interviewees how to achieve benefits and overcome challenges, their recommendations are also woven into this chapter. Although these themes are separated, they are highly interrelated as will be discussed in the next chapter.

## 4.1 The Perceived Definition of AI

As we have extensively highlighted throughout this thesis, the definition of AI is a heavily disputed topic, which scholars, industry experts, and business managers cannot seem to agree upon. During our interview process, we have observed a similar phenomenon. Out of the interviews we have conducted, none of the participants have the exact same definition of AI, but most of them converge around practical AI usage through machine learning, that AI is about replicating human behavior and automating human tasks.

Naturally, we have observed that AI developers on a general basis had a more comprehensive and advanced definition compared to business leaders and operators. For instance, one of the interviewees explained AI as something that “encapsulates information and knowledge. So, AI to me is the digitally usable condensate of knowledge so that humans typically generate more time.” (Interviewee 6, Case 4). Another pointed out that the definition of AI changes with the times, and that “AI is everything we cannot do yet” (Interviewee 10, Case 7). Nevertheless, most participants, even data scientists, expressed a struggle to define AI, and some of the answers were excessively comprehensive or linked to the practical use case.

The participants that did not directly work with AI development generally had more problems defining AI, and exclusively linked their answers to the practical use case. As one interviewee said: “I’m not sure about that.... I do not know if we should define AI, at least I do not want to.” (Interviewee 19, Case 11). After this, the person continued to explain around the question, and pointed out the functionality of the ML algorithms and how they thought it was a good thing.

The misconceptions and heightened expectations fueled by the hype surrounding the latter, and the mistaken belief that “AI is a form of magic” (Interviewee 2, Case 2), serve as an indication of a significant knowledge gap. For instance, one interviewee mentioned the need to label everything with AI as it is the “hot new topic” (Interviewee 14, Case 9), whilst another connected AI to intelligence and being able to learn by itself, highlighting that “AI sort of becomes a nirvana,

which is hard to achieve” (Interviewee 2, Case 2). They mostly connected weak AI with simpler statistical methods, ML, and regression analysis, and emphasize that ML algorithms have been around for a long time and are not as magical as the media portray it as. The latter was especially typical for people that had worked directly with AI. One of the participants even heavily contradicted our intelligence-related definition of AI, saying it was: “totally bull\*\*\*\*” (Interviewee 4, Case 3), before explaining that machines are not intelligent today.

Also, this signals a need to cool down the AI hype and be more realistic with the current state of the technologies. Here, it becomes essential to demystify AI, distinguishing between strong AI and the more traditionally applied narrow AI. The inherent uncertainty and lack of a clear definition of AI, which is intrinsic to the adoption of new technology, need to be addressed when incorporating AI into an organizational context. This is because the lack of a shared definition indicates that there is a lack of a shared understanding, which is a serious threat to the successful implementation of AI in organizations.

## 4.2 Attitudes and Expectations toward AI

Our interviews suggest that employees generally maintain a positive attitude towards the development and implementation of AI in their organizations. As anticipated, the developers of AI exhibit the greatest openness and enthusiasm to the implementation of the technology, whereas operators remain a bit more skeptical. Although the operators value the potential benefits of AI in streamlining their tasks, such as monitoring critical equipment, issuing notifications when required, and automating repetitive tasks, the initial impression of the technology is characterized by an uncertainty about how it works.

[I was skeptical at first] because it seemed so distant. It is so difficult to understand that a machine... can produce a result that ultimately turns out to align with what we are looking for. So, from being skeptical and finding it a bit strange, we have been reassured over time as we have seen some examples and witnessed how it works. (Interviewee 18, Case 11).

Contrary to what the media suggests that AI will take over jobs, this is not a common fear in our study. As one participant mentioned: “AI is not going to take any jobs, people using AI will.” (Interviewee 2, Case 2). As such, our findings suggest that the underlying fear stems not from AI itself, but from a fear of not being able to learn and efficiently utilize these new tools, potentially leading to job displacement. In other words, the underlying concern seems to revolve around being replaced with more technologically capable and educated humans. Nevertheless, several participants mentioned there was some skepticism and fear in the start of introducing AI into their company. A good example of which can be found in one of our cases, in which AI was removed from the internal project name as it generated fear within the company.

To ease employees’ fears and skepticism, some proposed to clearly communicate the methods and results of AI models to ensure transparency and trust. The communication should put emphasis on what possibilities the AI model will uncover, such as steering the employees’ focus towards higher value creation tasks, and not that it is going to replace their jobs, which will result in resistance and fear. However, this may be quite challenging as AI algorithms can become black boxes, which will be further analyzed in 4.4.2. *Trust in AI*. Although there can be some resistance, several interviewees emphasized that AI has come to stay, and as one said:

It is better to utilize AI than not to use it. That will come anyway. So, if you try to live in the dungeon and pretend that AI is not coming, you will be on the losing end of the business. (Interviewee 4, Case 3).

Over time, we find a normalization of AI tools in the work environment, which was for instance compared to once groundbreaking tools such as the calculator. “AI stops being AI as soon as you have it. Then the novelty wears off and the fancy machine learning product is becoming a calculator... and just another tool.” (Interviewee 12, Case 9). It is likely that the observed fear partly stems from job insecurities and uncertainty around the new technologies, which enhances the importance of gaining common knowledge of AI as early as possible to ensure

that workers have a more accurate view of their future job environment and security.

This leads us to the importance of expectation management. In case of high expectations, not only from the interviewees themselves, who emphasize that the same goes for clients and other organizational members, AI tends to disappoint due to the disparity between anticipation and reality. This is related to the findings from subchapter 4.1, in which participants realize that AI is not the magical strong AI, but rather narrow AI. This suggests that organizations should, from the early stages of a project, be careful in how they articulate and communicate AI projects, in addition to providing necessary training, as this term can lead to unrealistic expectations and negative emotions. By emphasizing the realistic capabilities and limitations of AI, organizations can help set appropriate expectations, forming a solid basis for building trust between employees and AI tools. Consequently, AI should not be perceived as a threat, but rather as an opportunity to augment productivity and job performance.

Many have very high expectations and believe that it will solve all problems. It's a bit like sprinkling some AI on the problem and everything magically resolves itself. That's not how it works. So, I believe that the key to having a positive experience is to manage expectations in a good way. (Interviewee 2, Case 2).

## 4.3 The Benefits of AI in Organizational Decision-Making

### 4.3.1 Accuracy and Efficiency

Many of the AI solutions could handle a significantly larger volume of data and perform computations at speeds far surpassing human capabilities. “Humans just can’t do it. It would take too much time, too many resources, and be too expensive. We now have ways to do it automatically.” (Interviewee 10, Case 7). This not only allowed for expedited decision-making processes that would otherwise be too resource-intensive and costly, but it also allowed for incorporating more variables into the decision-making process that might not have

been present earlier due to human limitations. As these models were trained on a knowledge base that was significantly larger than what a human possesses, it allowed the incorporation of more knowledge behind the decision, which potentially could, because of the larger knowledge base, avoid biases that might distort a human's decision. "It is not just faster and more accurate. If done right, it can also cover much more ground right, in a way. So, it gets rid of our own biases and our own shortcomings regarding knowledge and decision-making." (Interviewee 14, Case 9). Another interviewee, conversely, pointed out that "[AI] models are trained on human input, so it should, in a way, be as rational as it was before, although you get rid of the human factors, like being tired and those types of things. (Interviewee 15, Case 10).

A central advantage that was frequently underscored by participants was the ability of the AI solutions to achieve a high degree of accuracy of around 90%, with one even reaching 97%. This means that the AI models in our cases usually are correct 9 out of 10 times, which is good. However, this depended on how the accuracy was measured. In the prediction cases one could directly measure accuracy based on predicted versus actual values, but in some other cases they used a relative accuracy estimate comparing the AI model to what humans would have done. In the former cases, one gets an objective accuracy measurement, but in the latter, the measurements can be subject to human biases and their need to not admit mistakes. As one interviewee mentioned: "Humans will never admit they are wrong, right. Basically, we can get 92 to 95% accuracy easily enough, but getting a human to admit they made a mistake is super hard." (Interviewee 11, case 8). This suggests that the perceived accuracy of the AI model can be wrongfully low and is somewhat higher. Nevertheless, employees were by some interviewees perceived to have about 10-15 % failure rate, corresponding to 90-85% accuracy, again emphasizing that the accuracy levels of AI models and humans in our cases are somewhat similar.

More specifically, we found two clear examples of heuristics in cases 2 and 6. In the former case, a mechanic had very high accuracy based on a "knocking on equipment" heuristic. Here, the mechanic could accurately identify the state of a machine and predict when maintenance would be needed, based on very limited



information. In the latter, a water level operator with ‘sticking a finger in the air’-heuristic had somewhat worse accuracy than the AI model. In this case, the operator responsible for water pumps would simply make an educated guess as to how much water would be needed in the future based on very limited information.

The person who is responsible for the pumps comes to work in the morning, puts the finger out of the window, has like a table of, so it is expert knowledge of course, but in essence they are taking their finger out of the window to feel the air. They know what kind of day it is, and they have historic data that they look up and then they just set it [water level] and then they adjust. (Interviewee 6, case 4).

Generally, 90% was perceived as an optimal accuracy level depending on the industry and problem characteristics. In the wastewater bacteria case, for example, there are usually 10 000 bacteria per sample, so if they “get the bulk” right, they can “go with the preponderance of evidence, just like the human does” (Interviewee 13, case 9). In other cases, one NLP model with accuracy as low as 75% was favorable over traditional methods that were as low as 45%, and another case in the water industry where over 60% was seen as good due to “natural uncertainty” (Interviewee 2, Case 2). Nevertheless, some cases could not provide accuracy estimates due to either time constraints or lack of information access. For the former, 6-8 weeks in a consultancy project was seen as too little to extensively work on the accuracy of the AI model. This shows that accuracy measures are usually built up over time and is not something that is realistic to extensively measure in shorter projects.

Common to most of the cases was the infeasibility to achieve 100% accuracy because of an unfavorable relationship with marginal costs and decreasing explainability. As one closes in on 90 % accuracy, increasing this by 1% may not be worth it as it requires a disproportionate number of resources and decreases explainability. However, this accuracy varied. What is interesting here is that despite the overall accuracy of some cases going down, companies still found the AI investment worthwhile as it enhanced the efficiency of the decision-making processes in other ways. An example here is from cases 3 and 4, in which there

were a significant number of activities that revolved around routine screening of transactions or inquiries, which could easily be automated with an AI model with lower accuracy than humans. Central to this was the person at the end of the decision-making process, which could always control the process and determine the optimal output regardless of what the AI came up with, acting as a safeguard. Adding to this is that several companies have added a threshold probability value. For instance, if a model was less certain than 80%, it would automatically send the case to manual handling. However, accuracy estimates were at about 90% compared to humans, which means that there is not much of a difference in accuracy.

And then we always have to think about the consequences that there might be because it is quite difficult to blame a machine or a robot, right? If you have transferred quite a substantial amount of money to the wrong accounts, then who is to blame, right? The [company] as a whole is responsible for the created mistake. So that is why we, yeah, we always try to minimize the risks... have a “second pair of eyes” (Interviewee 16, Case 10).

By reducing the time taken to arrive at a decision, AI models thus lower costs by conserving resources that would otherwise be consumed in labor-intensive and manual processes. This conservation of resources allows organizations to utilize their workforce for higher value creation jobs, not doing repetitive and mundane tasks, thereby enhancing overall operational efficiency. A good example of the increased speed is the bacteria detection case, in which they were able to avoid waiting a week to get samples back and forth from local state laboratories and instead used cheap and mobile equipment allowing for 30-second hands-on processing.

#### 4.3.2 Expanded Information Base and Sustainability

The adoption of AI solutions can also discover new possibilities by utilizing an expanded information base, pattern recognition and prediction. This recognition ability combined with more relevant information enable the identification of

patterns that could be challenging or even impossible for humans to identify. “We understand cause and effect well and we can handle 2 things that vary but cannot handle 3-4 things. That is when the machine learning algorithms come in. AI is better at analyzing patterns.” (Interviewee 17, Case 11). Extending this finding, several participants highlighted the value of digitizing previously manual processes, particularly when combined with other relevant data. Thus, they can make more informed decisions, resulting in more rational decision-making processes.

I think that is the benefit, that it enables us to consider so much more, like more options. And then also more effects that are interconnected.... We can estimate other aspects beyond just the traffic flows... Traditionally, it would be rather difficult to holistically evaluate these multiple criteria at once. (Interviewee 5, Case 3).

By achieving more accurate predictions, organizations can better adjust their resource allocations, thus improving environmental sustainability. An example is from the Water Level Management case, in which more accurate predictions of water demand allowed for better supply adjustment, resulting in less carbon emissions and less waste. This was also central in other cases where incorporation of AI allowed for a more comprehensive information base, especially reducing transport emissions.

They are using this same modeling... to basically look at the emissions from the traffic and basically link this to the carbon neutrality target of the city... Then they see what activities they should take that have the biggest impact in the reduction of the carbon footprint (Interviewee 3, Case 3).

Common to several cases was the usage of existing AI algorithms. For instance, if there is limited data or speed concerns, transfer learning techniques can be used to mitigate it. There are often existing NLP and ML algorithms that can be applied and modified to a diverse set of problems. This helps developers as they do not have to do all the programming from scratch, again enhancing sustainability due to re-usage. A good example of this is the offshore wind case, in which there were

two datasets from an offshore wind farm. The first included comprehensive data gathered over a long period of time, whilst the other was very limited. By applying transfer learning techniques, they were able to infer from instrumenting 10-20% of windmills, to the remaining 80-90% with an accuracy of 90%. This proved that data limitations can be bypassed without having drastic impacts on accuracy. Also, several other participants highlighted the increasing use and availability of AI algorithm libraries, which have made it more accessible, easier, and faster to develop AI solutions.

### 4.3.3 Increased Adaptability, Reliability and Preservation of Knowledge

Our interviews also highlighted the digital flexibility in AI-based solutions. The AI models are relatively easy to retrain or adapt to either changing business environments or identified model deficiencies, thus being able to provide an advantage in business agility. This is particularly notable when compared to a human's restricted adaptability in circumstances like a sudden change in tasks or the environment. NLP models are also able to handle unstructured data, decreasing the dependence on large amounts of structured data which is hard to come by.

As AI is inherently digital, you can always change it. And that is also an advantage in a business sense, right? If you have 50 experts sitting in a room just doing analysis, and tomorrow no clients want this analysis anymore, you do not know what to do with those 50 experts.... It's much easier to write a couple of lines of code and retrain a model than it is to retrain a person or our entire organization. (Interviewee 13, Case 9).

In addition, AI tools are a consistently available resource that is operating at optimal capacity around the clock. This is a significant contrast to human operators whose productivity and attention span often fluctuates throughout the day and require regular periods of rest. As such, the inherent reliability of AI tools provides an edge in enhancing operational consistency in an organization. This

would be especially important in organizations that rely on one or a few persons to effectively address an issue:

There are not a ton of advantages over the traditional method other than the AI does not get sick and does not want a day off. But that is a relatively good advantage. Also, it does not quit and go somewhere else.... Most sites have a biologist, so if that guy is out, then they do not do it that month or that week. If he decides to fall off a ladder or something, suddenly you are not using this analysis tool anymore. (Interviewee 13, Case 9).

Further, by incorporating AI solutions across different offices, more reliable decisions could be taken across organizational departments. As evident from several cases, humans make mistakes, but AI could be used to minimize the variations in their performance.

Another thing is that there is a lot of bias in machine learning and in it as a concept. But it is also quite clear that there is a lot of bias on the part of supervisors. Our insight work shows that there is a very different process from office to office with how they assess candidates... So, if you put it in a system, the idea is also that you might want to get a slightly more equal assessment in all offices throughout [the country]. (Interviewee 8, Case 6).

In addition, AI can preserve knowledge. This was a critical issue in the bacteria detection case as most of the researchers in the field were soon to retire, introducing the risk of losing valuable knowledge. Therefore, it was important to digitize knowledge rapidly, which was done by incorporating expert knowledge into the training process of the detection algorithm. This was also a reason as to why some organizations chose to train their algorithms with supervised learning to utilize expert knowledge.

## 4.4 The Challenges of AI in Organizational Decision-Making

### 4.4.1 Problem Identification and Comprehension

Most participants highlight the importance of the starting point in the AI-driven decision-making process. It was regularly mentioned that many organizations start to experiment with AI without trying to identify and solve existing problems. This was particularly connected to the hype around AI and to an initial excitement around experimenting with the newest and most fancy technology. However, by not solving an existing problem, it is very unlikely that the AI model will yield any value to the organization or client. Thus, organizations risk investing a significant number of resources into projects that will never make it out of the research & development (R&D) department and is ultimately worthless, apart from potential learning effects. For instance, one case participant mentioned numerous “wasted” projects due to this challenge. To solve this, several participants recommended to be more realistic when it comes to AI, but still to be bold and try things out. For instance, the “boring” applications are usually the most impactful.

So, we go and talk to users first and then we work with the AI algorithm. That has worked very well for us... whenever we develop something, people are already there using it for us. We have made 10 AI algorithms and all of them have been used by the organization. So, our success rate is 100%. (Interviewee 1, Case 1).

However, the organization also needs to understand the problem. As one interviewee mentioned: “Are you really sure you understand the problem that you are trying to solve is like the first question. And typically, the answer is no, and I think that is the 80% [of the failure rate] right there.” Breaking the problem down into smaller segments was mentioned as a helpful way of solving this, which can lead to a realization that AI is not the best solution to the problem. Case participants also favor a more gradual approach to developing AI, asserting that it is important to restrict the AI usage to an understandable level in the start, gradually increasing the difficulty. Understanding a problem and pursuing gradual

development would, however, be highly inefficient without starting with an identified problem. Nevertheless, these are processes that interviewees mentioned would take a lot of time, but recommended patience and trusting the process as keys to succeeding.

Our analysis suggests that neglecting problem identification, understanding and a gradual development approach can be symptoms of the AI frenzy, which shows that the current environment can lead to ignoring sound decision-making. Nevertheless, these are only some of the many prerequisites to overcome AI challenges. Trust is another one of these factors.

#### 4.4.2 Trust in AI

Trust in AI is critical for AI to augment organizational decision-making processes, which depends on explainability, biases and control. This is a central part of bridging the gap between AI and potential users. We found the explainability factor is contingent on the nature of the use case, which is essential for uncovering biased or technically faulty algorithms and achieving transparent and controlled decision-making processes.

As the accuracy of an AI model is inherently linked to how complex it is, it creates a transparency issue. As a model increases in accuracy, it inevitably becomes more complex and hence more challenging for humans to comprehend, a so-called “black box”. A significant number of respondents underscored this as one of the major disadvantages of using AI in organizations, as it would not be ethically sound to blindly trust and act on the outputs without understanding the underlying processes. As such, the response to how much each respondent sees this lack of transparency as an issue in their application varies greatly, where organizations without direct human implications tend to prioritize accuracy over transparency. In contrast, organizations with direct human implications favor a more balanced approach to accuracy and transparency. An example from the Bid Robot case shows that black boxes can lead to the AI development being shut down:

We collected a large amount of data from payroll systems, project systems, customer evaluations, and so on, and trained the model on this. Then we asked the model, 'Which employees are likely to resign within the next three months?' We achieved an accuracy rate well above 95%. We attempted to understand the underlying mechanisms of this model, but we could not.... As it did not align with our values, we decided to shut down the project. (Interviewee 11, Case 8).

Although explainability might be most desired in some contexts, the resulting decrease in accuracy can decrease the trust towards the AI tool. This is because the AI model would more often fail to produce the desired results, which would lead to the AI being less dependable and trustworthy. Therefore, this is an effect that can reduce the value of increasing explainability, hence decreasing accuracy, in the first place. Although we recognize that some degree of explainability is crucial for the development of trust, we propose that one should acknowledge that there seems to be an inherent tradeoff between accuracy and explainability. With current technology, an emphasis on explainability will limit accuracy, and as such, limit the potential value of developing and implementing AI in organizations in the first place.

Furthermore, several participants highlighted potential biases in AI algorithms as one of the most significant drawbacks of using AI in organizations. This issue becomes particularly important when models are trained on historical data involving humans. Unfortunately, as human history includes countless instances of injustice and unfairness, there is a risk that these biases are encoded in a model, as algorithms can perpetuate or amplify existing societal biases. These concerns are particularly highlighted in automation tasks with limited human intervention, as relying on AI tools can provide a sense of loss of control. Consequently, organizations may experience a diminishing sense of agency and decision-making authority when AI algorithms manage vast data quantities, without human interference. "Bias is the most important [disadvantage]. In a way, you're giving away control, right? If a human, for example, would read thousands of comments, they might come to a different conclusion than the model." (Interviewee 10, Case 7).



Even though having a human at the end of the organizational decision-making process is central for control, it is imperative to make the operator trust the output of the AI model. In most of the cases we have examined, there is a human that acts on the model's output. As a result, an AI tool will never be beneficial for an organization unless the operator trusts the model's output and understands how to act upon it.

AI can make a decision-making process more rational, but ultimately, the irrational decision-making lies with the human at the end of the process.... He must truly believe that what the model says is correct. If he has experienced too many false alarms, that's where the problem lies. It is not that the model is wrong, but rather that he does not bother getting out of the car for yet another false alarm. A model will do what it has been trained to do, but what you are actually trying to achieve happens out in the field with people. (Interviewee 2, Case 2).

#### 4.4.3 Costs and Complexity of Data Management

Several respondents highlighted the cost associated with developing and implementing AI models as a significant disadvantage, where the financial burden of setting up and maintaining AI systems can prove prohibitive for many organizations, especially smaller organizations, and start-ups with limited resources. The implementation process itself is perceived as costly and intricate, which involves training leaders and employees tasked with interpreting AI outputs and ensuring that the organization's internal procedures are equipped to utilize the new technology. However, some of the cases were from 2019 and the AI evolution has come a long way since then, leading to better and faster AI development.

A significant challenge shared by most of the respondents was the limited availability of data. Given the nature of AI models, which require significant amounts of high-quality, accurate, and up-to-date data to train and validate their performance. Obtaining access to training data was underscored as a significant barrier, particularly for smaller organizations and start-ups, not only due to the

inherent costs of data acquisition but also the substantial time investment and necessary knowledge required to format the raw data into usable data.

In cooking, it's really important to get high-quality ingredients and prepare them the correct way... so that you cut them in the right shapes, sizes, and so on. Do the right preparation before you even start doing anything.

[Developing] AI is no different. Never underestimate how long it takes to gather and prepare the data you are going to use (Interviewee 10, Case 7)

The importance of managing data seems to be dependent on the use case itself. Naturally, a greater level of importance to data bias is to be found in the cases directly impacting humans. Several of the participants have mentioned limiting the use of personal details to a bare minimum, in addition to exploring the use of synthetic data as potential solutions. In addition, the acquisition of data was frequently tied to privacy considerations, which often impose constraints on the type of data that could be collected and used. As one interviewee emphasized: "The biggest [challenge] that I have seen is data privacy and how to deal with it. If we get data that is somehow raw, we need to be really strict on developing anonymization." (Interviewee 4, Case 3).

Further, there is a scarcity of high-quality and up-to-date data, which significantly influences the accuracy of AI models. This results in a gap between the potential and actual performance. In other words, with inaccurate or wrong data it is likely to end up with an unsatisfactory result. This discrepancy was identified to be a major obstacle to securing continued managerial support, both financially and in terms of time allocation during a project's developmental phase. "Even if you have the best machine learning models, it is still going to be a garbage in, garbage out-situation. Everything depends on the quality of the data." (Interviewee 2, Case 2).

What exacerbates these issues is the inherent uncertainty in the development process. When starting with developing AI models, it is hard to tell how many hours will be required to achieve the necessary accuracy. As one interviewee mentioned:

But when it comes to machine learning, we cannot tell if we invest 100 hours, then we will find this pattern or solve this with the algorithm. Or if we invest 1000 hours, we still do not know whether we will be left with a functional model or not. (Interviewee 11, Case 8).

In addition to data challenges, several participants highlight the importance of having access to the right talent throughout the process of developing and implementing AI. In the development process, for example, there is a greater need for someone who can gather, provide, and clean data. In the implementation process, on the other hand, there is a greater need for someone who can take the AI model and successfully implement it within the organization and educate the employees so that it provides the intended value. This necessitates a mix of individuals with diverse backgrounds, including those with technical expertise as well as non-technical professionals. Revisiting the cooking metaphor once more, one of our interviewees described it this way:

[When the preparation is done and] you start cooking... you also need chefs that have made the dish before, and preferably with different specializations and skills. When the dish or product is done, you need good waiters to deliver it, right? There are a lot of steps [to developing and implementing AI tools], and every step is important and requires different people with the right skills and resources. (Interviewee 10, Case 7)

#### 4.4.4 Regulatory and Economic Uncertainty

Firstly, participants expressed concern over existing and outdated laws and potential future legislation that could restrict AI usage. This concern is especially prominent in regions such as in the EU, which has robust data protection regulations such as GDPR, where uncertainty around future regulations acts as a deterrent when considering implementing AI tools. As one interviewee mentioned: "We are working with laws that were developed in the 90's that could not foresee how rapid the data development would be at all. So, it is a tremendous challenge." (Interviewee 15, Case 10)

Further, regardless of communication with the legal department, the CV & Job Matching case was shut down partly due to regulatory uncertainty and risks within AI in recruitment. This also presented some difficulties as they were not allowed to train the NLP model on personal data but was solved by just comparing text and training the model on news articles instead. This signals that despite all the measures one can take to realize benefits and evade the other challenges, regulations can be detrimental to utilizing AI in organizations.

Secondly, economic aspects can also present a challenge regardless of how far along the project is. This was again prevalent in the CV & Job Matching case, in which they faced budget reductions and were thus forced to alter their resource allocations. Adding to this was the fact that the AI project was an internal process improvement of which functionality could be somewhat covered by existing acquired solutions. Some other cases also had this issue as AI development was characterized as research and development (R&D) activity that would be less prioritized when at conflict with core business tasks. Interestingly, one interviewee hoped that if they automated enough of their consulting projects in the future, they could focus more on R&D.

Often, it [a process innovation] is perhaps in areas where you have something that works today, but which could be improved. [The company name] like many others now, must cut their budgets. Then it is easier, I think, for the management to cut new development, instead of what already exists, in a way. It is a bit more like experimenting. (Interviewee 8, case 6).

This highlights how important it is to spend enough time assessing risks and uncertainties before investing in AI projects. Failing to address these can lead to wasting valuable organizational resources, although we recognize that these factors are highly exogenous. Adding to this is the organizational complexity and inertia hindering effective AI development.

#### 4.4.5 Organizational Complexity and Inertial Forces

The organizational complexity, especially in terms of inertial procedures, was a recurrent problem throughout our interviews. This complexity was particularly severe in large and established organizations where outdated structures and procedures often proved inadequate to meet the required pace of current technological advancements. Smaller organizations, however, were to a larger degree more capable of implementing AI tools in a shorter time frame. Several of the cases were also the first AI projects in their companies, meaning that they for instance lacked successful projects they could look to for help and organizational strategies and routines for collaboration around new AI projects. Further, securing approval for internal testing and implementation processes was described as demanding and time-consuming, often resulting in considerable delays or even project cancellations.

There is a very complex application landscape in such a large company. Everything needs to fit together, and everything has to go through numerous gateways before anything can be put into production. It takes an extremely long time and is challenging to work with. Out of the six months this project lasted, three months were actual work, and the rest involved obtaining approvals. (Interviewee 15, Case 10).

Furthermore, several participants highlighted the organizational infrastructure and its supporting role in the development and implementation of AI as equally critical for success, and just as important as the model itself. As such, introducing the importance of developing an appropriate organizational structure that is fit to handle both the development and implementation processes of AI. Numerous participants highlighted the importance of incorporating AI into a company's overall strategic framework. As many of an organization's objectives and goals are optimistic in nature, they might be unattainable without incorporating machine learning and artificial intelligence. “[Developing AI before the organization is ready] is typically where organizations burn a lot of money. If the processes and the organization are not ready for AI, then you should fix that before developing AI [solutions].” (Interviewee 12, Case 9).

However, organizational inertia can be highly dependent on their markets. Several participants from the consulting industry connected this to a current transition phase in their markets. This is not only due to conservative customers slowly adapting to the new technologies, but also that consultancies' business models are based off being paid by human working hours, and not machines. Naturally, this business model does not strongly incentivize efficiency gains as more human hours will yield more revenue. Although enhancing efficiency can for instance allow consultancies to reduce their prices in tender competitions, there seems to be strong inertial forces in place hindering this, both within the organization and the clients. One interviewee made an example of a market that had been the same for 50 years, not willing to change and adapt to AI no matter what, whilst another mentioned that:

I just think you need to read the market and you need to dictate your choices based on what the market would accept or not. At the moment the market is not totally accepting of it. It is a transition period, and I am sure once the market wants to accept it, [the company] will start wanting to do it, right. But you cannot just force things on clients. (Interviewee 7, Case 5).

Although the implementation process in larger organizations can be time consuming, this is not exclusively negative. Several interviewees mentioned that this allows for a thorough review of how an AI tool is built, and if the ethical and privacy implications of the tool is in line with current regulations and internal goals and routines. One interviewee further state that the inclusion of clear regulations and ethical standards could increase the level of trust towards an AI tool:

We have very strong regulations.... You need people to trust this [AI], and one way that you can make people trust it is by either adhering to some law or regulation that everyone agrees on, and everyone trusts, or by proving very, very transparently, deliberately, and painstakingly that you are to be trusted. (Interviewee 12, Case 9).

## 5.0 Discussion

This section aims to discuss our main contributions to theory and managerial practice considering existing literature. We contribute with 4 main themes: 1. Rationality and Accuracy, 2. Trust in AI, 3. Organizational Structure and Strategic Goals, and 4. Problem Comprehension. In addition, we will discuss possible limitations of the study and identify areas of further research.

### 5.1 Theoretical Implications

Even though most of our findings converge with existing literature, we also contribute to gaps in the literature. More specifically, we provide new insights into how the context of use cases can affect accuracy and explainability through assessing whether humans can be negatively affected; how the AI-based decision-making structures of Shrestha et al. (2019) can be modified to represent the complexities of AI augmenting humans more accurately, and overall, how the different advantages and challenges are tightly interrelated and will depend on knowledge about AI.

#### 5.1.1 Rationality and Accuracy

Our study suggests that AI can augment decision-making processes, but that this is contingent on numerous factors. However, this improvement cannot be exclusively attributed to accuracy. For instance, mechanics can, by knocking on some equipment or ‘sticking a finger in the air’, gather limited and relevant information, and propose a suitable decision to a medium to high degree of accuracy. Given that AI algorithms can reach a similar degree of accuracy, which our results support, the human is often able to perform similarly rational decision-making processes, all else equal. This is in line with the findings of Tversky & Kahneman (1974) and Gigerenzer & Gaissmaier (2011) that heuristics can be a useful decision-making strategy in some cases characterized by uncertainty. On the one hand, this alone would not justify the investment in AI as the organization would likely be better off just keeping the human. On the other hand, combining speed, reliability, and sustainability with similar levels of accuracy tends to justify

the investments. For example, an AI system is far more reliable than a mechanic, given that it is always at work, does not take sick days, gets distracted, “or decides to fall off a ladder” (Interviewee 13, Case 9).

Another aspect that can justify AI investments is that these models are closer to perfect rationality (Simon, 1979) than what humans might be. They can, through digital networks and the larger processing power of data, account for more information on more alternatives than a human ever could (Shrestha et al., 2019). For example, by connecting a planning process to carbon emission data, it is possible to account for relevant perspectives that were not considered before and thus reach higher predictive accuracies. This shows a transition from humans doing the information search of decision-making processes before, characterized by bounded rationality, to a search process closer to perfect rationality (Simon, 1979). Although this can initially lead to a belief that AI will augment the organizational decision-making process through increased accuracy, only the Transport Modelling case supports this. The scarcity of support can be attributed to the lack of sufficient high-quality training data and that organizations are using simpler AI models to ensure explainability, which could be affecting accuracy negatively. Thus, incorporating more information and parameters into AI-based decision-making can lead to augmentation of humans, but it often requires that the challenges are addressed. However, this is easier said than done as our analysis shows. This suggests further support for Gigerenzer and Gaissmaier’s (2011) work on the value of heuristics compared to comprehensive AI analysis, as the latter can be highly challenging.

The specificity of decision search space and interpretability of more complex AI models (Shrestha et al., 2019) can be more effectively addressed by humans than AI as their intelligence is not bound to a particular use case and that humans can increase overall explainability that AI can lack. Regardless, we found that the size of the alternative set was the most important in that AI allowed for processing of vast amounts of data that humans were unable to tackle, significantly increasing the efficiency of the process. Further, we found little support for AI’s significance through enhancement of decision-making speed and issues of replicability (Shrestha et al., 2019). Overall, the organizations in our study aligned their AI-



based decision-making structures with “AI to human sequential decision making” (p. 6), effectively combining the strengths of AI and humans, and not amplifying each other’s weaknesses. However, whether high interpretability is achieved in this structure will depend on how advanced the AI technology is, which Shrestha et al. (2019) fail to comprehensively address. For instance, several of our cases had low interpretability although there were humans involved at the end.

When discussing the role of AI in organizational decision-making, emphasizing the role of human judgment is crucial. Nearly all of our participants have highlighted that a human operator is a part of the process, and as such, able to oversee the process in case of errors. This is further highlighting that the need for human oversight in AI processes is crucial for avoiding potential mishaps, which is in line with the upcoming EU AI Act (European Parliament, 2023): “AI systems should be overseen by people, rather than automation, to prevent harmful outcomes.” Also, as the cases have seen a threshold value that sends lower confidence results to manual handling, this can lead to improved accuracy and fairness of the systems, consistent with Shrestha et al. (2019).

However, as humans can be irrational, for instance due to heavily biased feelings (Slovic et al., 2006), it is important to stress a factor that could limit the accuracy of an AI model: the operator acting on the output of the model. An operator often makes the final decision on whether to act on the output a model has given or not, meaning that it does not matter how accurate or rational the output is if the operator chooses not to act on it. For instance, the operator could be stuck in old routines and be afraid of the new technologies, rendering the processing of the AI model useless. The ability to bridge AI and users would naturally depend on trusting the algorithms.

### 5.1.2 Trust in AI

Scholars have emphasized trust as a critical factor in AI implementation (Kolbjørnsrud et al., 2017; Kolbjørnsrud, 2023), which is in line with our findings. In addition to this, we have observed a struggle with the inherent tradeoff between accuracy and transparency, supporting the findings of

Eschenbach & Warren (2021). As we have observed, these two aspects have a direct impact on the level of trust towards an AI tool, and therefore, the probability of a successful implementation. As such, the findings of this study reinforce the critical role that attitudes and trust play toward AI in the successful integration of AI into organizational decision-making processes.

#### 5.1.2.1 Managing Trust in Organizations

Consistent with prior studies (Castelvecchi, 2016; Kolbjørnsrud et al., 2017), our investigation suggests a direct relationship between individuals' attitude and their level of trust in AI with their knowledge and hands-on experience with AI tools. Building on the findings Jacovi et al. (2021), there are two types of trust towards AI tools, namely intrinsic and extrinsic trust. As such, trust towards in an AI tool can be obtained through explanation, or in other words, transparency, or through the evaluation data, or in other words, consistently accurate outcomes.

##### *The Role of Intrinsic Trust*

According to Jacovi et al. (2021), a user will gain intrinsic trust towards an AI tool when 1. The user successfully comprehends the true reasoning process of the model, and 2. The reasoning process of the model matches the user's priors of agreeable reasoning. Due to the inherent trade-off between accuracy and explainability, as illustrated in our findings, this aspect of trust is difficult to manage, as increased explainability can reduce the overall accuracy of the model. Providing new theoretical insights, we found this to be dependent on whether the cases had direct human implications, in which organizations with direct human implications tend to favor transparency over accuracy. However, the AI Act (European Parliament, 2023) can alter this in collectively pursuing more explainable models, but with adaptations to the risk level of the use case. Regardless, it could be beneficial to consider other options to increase or replace the need for transparency, one of which could be legislation and strict organizational routines. An interviewee compared this to the construction and use of bridges, where most people are comfortable walking over a bridge, knowing that it is built under the adherence of strict rules and requirements, thus making transparency somewhat redundant.

The second aspect is that the reasoning process of the model matches the user's prior knowledge of agreeable reasoning (Jacovi et al., 2021). This is consistent with our findings, which suggest that practical experience and education can provide a more accurate and realistic view of AI capabilities, which increases the level of trust. It is, however, important to note that not every employee requires comprehensive technical knowledge to trust AI. As Jacovi et al. (2021) states, it is not a prerequisite for trust that a user understands the outcome, as long as the user believes that the outcome is correct and trustworthy. Like our trust in common devices such as computers and televisions, the primary focus lies in their functionality and reliability. Most employees possess a rudimentary understanding of how these devices operate. However, the specifics of their inner workings and assembly would render most employees clueless. For instance, most employees trust Microsoft Excel to correctly execute multiplication despite not fully grasping the underlying computations and how the product is built. This same perspective should extend to AI tools; the emphasis should be on their functionality and reliability, rather than on intricate technical details. Ensuring employees have a basic comprehension of AI can help alleviate fears and increase acceptance.

#### *The Role of Extrinsic Trust*

To obtain trust, one approach is with Extrinsic trust, or in other words, consistency. As such, the developers need to prove that the AI tool is consistently accurate over time. According to Jacovi et al. (2021), there are two methods of obtaining extrinsic trust: 1. By proxy, and 2. By observation.

First, according to Jacovi et al. (2021), a user can obtain trust towards an AI tool by being influenced by an expert human opinion. By being assured that the outcome of the AI tool is correct and trustworthy, a general level of trust can be achieved. This is consistent to a method used in our cases, where experts presented the model to the operators.

Second, a user can obtain trust towards an AI tool by observing that a model produces consistently accurate outcome over time, in different specific contexts (Jacovi et al., 2021). This is consistent with our findings, which suggest that the

level of trust towards an AI tool is directly tied to the operators experience with the tool. The authors further exemplify this by showing that if a model performs equally for two large collections of people of different races, a user may deem the model more trustworthy (Jacovi et al., 2021). This, however, introduces a repeatedly highlighted concern in our study, namely the significant challenges of obtaining and cleaning the data used to train an AI model. If the datasets contain biases toward race or gender, these biases could inadvertently be learned and amplified by the model (Bolukbasi et al., 2016). This is further consistent with what we have observed in our study, where nearly all of the interviewees emphasized the importance of accurate, high-quality data and unbiased data management.

#### 5.1.2.2 Attitudes and Knowledge Towards AI

We have found that the knowledge an organizational member has about AI is decisive for developing trust in the technologies. As evident in the current situation, there generally is limited knowledge about AI in organizations (Scarpetta, 2023). Combining limited knowledge with potential fears and biases from media can be a perilous combination (Nader et al., 2022). If organizational members are mainly exposed to and believe in entertainment media's dystopic portrayal of AI and the fact that many top organizations have requested a halt in the technological development, signaling a control loss of the groundbreaking technology (Kahn, 2023), it is likely that they will experience negative emotions toward AI, such as fear and skepticism (Nader et al., 2022). This enhances the importance of raising organizational awareness about AI, which could also result in enhanced intrinsic motivation towards using the new technologies, if coupled with autonomy (Deci & Ryan, 1985).

It is evident from existing literature and results that AI has the potential to evoke adverse emotions among organizational members through fear, skepticism, and/or disappointment. Even though the AI label is avoided in communications and projects, there can still be resistance and adverse emotions due to AI changing jobs (Leonard-Barton & Kraus, 1985). Heuristics can be especially alluring in such complex situations, in which humans can rely on the likes of the availability

and affect heuristics to ease their efforts (Slovic et al., 2006). However, “The worst thing a manager can do is to shrug such resistance aside...” (Leonard-Barton & Kraus, 1985), which can be avoided by addressing the need for relatedness and increasing intrinsic motivation (Deci & Ryan, 1985). This suggests that managers need to address the feelings of their employees to provide a sense of security, no matter whether they are overly positive or negative. By doing this, they can create a better climate for developing trust toward AI.

In conclusion, our study aligns with the findings of Haefner et al. (2021), that designing an AI tool that humans can interact with and adequately trust are important challenges to overcome for successfully developing and implementing AI in an organization.

### 5.1.3 Organizational Structure and Strategic Goals

The fact that structural inertia is a result of an organization’s age and size is a well-known phenomenon (Hannan & Freeman, 1984). In this study, we have found equivalent results, where younger and smaller organizations seem to be able to implement AI tools more rapidly as a part of their decision-making process, compared to their counterparts. According to March (1991), the innovative capability of an organization is inherently dependent on the balance between exploitation and exploration, representing a common issue we found of balancing core business functions and innovation. Further, we have observed that a lack of a clear company-wide vision about the technological future and how the organization aims to utilize AI has been an internal obstacle when developing and implementing AI. As such, in line with the findings of Adams et al. (2006), we believe that incorporating AI into the organization’s strategy, for instance in the form of innovation goals (Yukl & Lepsinger, 2006), can play a crucial role in shaping an organization's approach to successfully develop and implement AI. By effectively implementing such a strategy, organizational members can be allocated more time and resources to pursue new learning and innovation. As Yukl & Lepsinger (2006) emphasized: “When there is a specific innovation goal for which people will be held accountable, this mental activity is more likely to get the attention and effort it deserves.” (p. 5). As such, we believe that this can

improve interdisciplinary collaboration, as especially prominent in the implementation of the Fleet Management case.

Regardless, effectively implementing AI strategies requires that leaders understand the technologies, which seems to be a common problem. Not being able to understand this can lead to managers failing to communicate transparently about the benefits and limitations of AI, not addressing employees' concerns effectively, and therefore not following the advice of Kaplan & Haenlein (2019).

Our findings suggest that consulting companies are slow to internally utilize AI in their decision-making processes, because their business model incentivizes billable hours over efficiency improvements. This is consistent with the findings of Crisan and Stanca (2021), who found that only a small portion of consulting firms have replaced their traditional business models with those based on new digital approaches and technological innovations, urging the need to change business models to incentivize technological adoption. Furthermore, we have found that internal digital transformation in consulting companies is dependent on the market and whether the market is accepting new innovations. This is again consistent with Crisan and Stanca (2021), who uncovered that the digital transformation of consulting companies is closely connected with their external triggers, or in other words, the demand in the market. We have, however, found that in our cases, several interviewees recognize the need to change, and that the market currently is in a transition period converging towards AI acceptance.

Extending this, these significant inertial forces signal strong path dependencies that can be exceptionally hard to break (Arthur, 1994; Cyert & March, 1963; Sydow et al., 2009). However, dissolving such path tendencies is possible, and can be a result of unforeseen exogenous forces, such as “shocks, catastrophes, or crises” that are likely to “shake the system” (Arthur, 1994; Sydow et al., 2009, p. 701). The recognition from the participants that they need to change their business models signal that these patterns are, in fact, starting to change. This also indicates that AI can act as an unforeseen exogenous force, which is reasonable due to the ground-breaking developments in the last years shown through the likes of ChatGPT.

#### 5.1.4 Problem Comprehension

We established that starting with a problem was essential in the decision-making process, as failing to address this can virtually render the project useless. Because implementing new technologies can be considered as major organizational change processes, change management principles can be of help here. Our findings converge with the likes of Kotter (1995) in that it is vital for successful change processes to have a “burning platform” and maintain organizational momentum around it. By starting with a problem, the AI change processes already have this platform in place. Therefore, they avoid some of the challenges connected to first developing AI algorithms and then trying to force technological change on employees. Central to Kotter’s (1995) view is the importance of cooperation and motivation amongst organizational members, which evidently applies to today’s AI change processes as well. Extending this, to force change on employees violates Ryan & Deci’s (1985) autonomy contingency of intrinsic motivation by “thwarting people’s innate psychological needs.” (p. 71), which is supported by the likes of Leonard-Barton & Kraus (1985).

As strategy and goals affect organizational decision-making processes (Cyert & March, 1963), problem comprehension also affects the information gathering and analysis of organizational members (Simon, 1947). By not having a solid problem comprehension, the following decision-making process can be error prone and biased. Conversely, if organizational members spend enough time to uncover root causes and a deep understanding of the problem, this can allow for consideration of more relevant alternatives, being able to evaluate the consequences more accurately and eventually choosing the more appropriate alternative. Starting with a problem and fully comprehending it converges with Simon (1979), in which a well-founded problem identification and comprehension can lay an important foundation for close to perfectly rational decision-making processes.

The finding of pursuing gradual AI development is highly consistent with Engelbart (1962). His strategy for research toward augmenting human intellectual effectiveness emphasizes the precepts “to pursue the quickest gains first, and use

the increased intellectual effectiveness thus derived to help pursue successive gains.” (p. 3). By starting slow and initially focusing on quick gains from implementing the “simpler” models of AI, such as machine learning algorithms replicating decision trees, this can allow organizational users to enhance this augmentation effect even further. Thus, after having learnt the easier models, it will be easier to progress into more complicated AI, such as deep learning.

The excessive hype and confusion around AI as found in the literature (Leonard-Barton & Kraus, 1985; Siegel, 2023), are mirrored in our findings and analysis. Interestingly, participants connect AI to AGI, and not to narrow AI and ML methods. As the AI label can be highly misleading, we believe in two solutions: 1. Address the knowledge gap and ensure more realistic expectations, which undoubtedly is a large change process that will take significant resources and time (Brynjolfsson & McAfee, 2014; Leonard-Barton & Kraus, 1985), or 2. Change the name of ML-based project to not include AI, a quick-fix, like in the CV and Job Matching case and as Siegel (2023) recommends. Pursuing both can be an effective way of managing expectations in the short-term, but also long-term in ensuring better and more accurate understandings of the new technologies. This can therefore help to cool down the AI hype and contribute to more business value through the more simple and practical use cases.

## 5.2 Managerial Implications

It is evident that using AI effectively in organizations is a demanding task, exacerbated by the large investments of time and resources required, in addition to the numerous pitfalls along the way. In our study, we have gained insights into the complexities of AI and what factors that can lead to either success or failure. In this section, we aim to provide a summary of the most useful managerial advice by combining our empirical findings and theoretical foundation. We believe that by focusing on these recommendations, organizations could increase the success rate of AI development and implementation, as earlier discussed in the introduction.



1. Carefully manage expectations in the organization regarding what AI is and what the intended AI tool is capable of.
2. Use AI to innovate based on existing organizational problems, not just to experiment with technology. This process can be hard, and it is normal to fail, but starting with AI as fast as possible is paramount to not being at a competitive advantage in the years to come.
3. Assess different relevant technologies before implementing AI to ensure that there are no other more suited, and possibly simpler methods to addressing the problem.
4. Assess whether the decision-making process can affect people negatively to ensure corresponding transparency.
5. Understand the problem at hand from a legal view and what the risks are of using AI in your domain. The EU AI Act can be a helpful starting point.
6. Assess whether you have access to enough quality data to effectively train an AI model and acknowledge that data gathering and cleaning are straining processes. If this is in place, start to experiment gradually.
7. Manage trust by enhancing the understandings of how the AI models works and ensure that they produce accurate results consistently.
8. Include humans at the end of decision-making processes with mechanisms to allow for higher safety, accuracy, and fairness.
9. Use change management principles to manage the change processes that AI brings into the organization. This can be to create a burning platform, motivate employees to support the new technology by enhancing their AI knowledge, autonomy, and relatedness, and to leverage successful use cases.
10. Be patient and acknowledge that the AI wave, like the sustainability wave, will take time due to industrial and organizational inertial forces. This can take a significant amount of time if AI is being utilized in the organization for the first time, as existing procedures and organizational designs might not fit to effectively implement AI.

### 5.3 Limitations

In this master's thesis, several limitations need to be acknowledged that might have impacted the outcomes and overall findings of the study. Firstly, our findings are based on a multiple case study design with cases from different industries. Although this method has enhanced the breadth of our findings, it also introduces an inherent level of heterogeneity that can complicate the identification of common patterns or trends (Ridder, 2017). The broad spectrum of industries studied, each with its unique contexts, cultures, and decision-making styles, might have shaped how AI was utilized and perceived in each use case. Furthermore, it should be noted that our insights are based on a specific set of case studies and therefore should not be uncritically applied to all organizations across all industries.

As the interviewees held different educational backgrounds and levels of experience, this could have introduced bias into our data. These differing backgrounds could result in varied interpretations and emphases on the role and impact of AI in decision-making processes. For instance, individuals with a technical background might emphasize the functional or performance aspects of AI, while business-oriented individuals might prioritize economic considerations, organizational culture, and soft skills. While this subjectivity is challenging to control for, we have attempted to mitigate this issue by balancing the number of interviewees with different backgrounds evenly. Although we believe this measure has been beneficial for a comprehensive understanding of each case, we recognize that this potentially could limit the cohesiveness of the findings drawn from the data.

As shown in Table 2, we were in some cases only able to conduct interviews with a single individual per organization. This leads to two potential limitations. The first raises concerns about response reliability, as our insights from these organizations heavily rely on a single individual's perspective and memory, which may not fully represent the organization's diverse experiences or viewpoints. Secondly, it may limit our in-depth understanding of how AI is embedded within

the organization and how it augments decision-making processes due to the lack of multiple internal perspectives.

Although qualitative data can provide rich and detailed accounts, our study may have been susceptible to interviewee bias, which is a prevalent issue in qualitative research where participants might either consciously or subconsciously present themselves or their organizations in a particular light (Qu & Dumay, 2011).

Despite our efforts to mitigate this bias through the careful phrasing of questions, creating an environment for open discussions, and analyzing and presenting the data in an objective manner, it is nearly impossible to eliminate this risk entirely.

Despite these limitations, we are confident that our study offers important insights into the augmenting role of ML and NLP in organizational decision-making processes. It is important to note, however, that the development and implementation of AI in organizations is a highly dynamic and rapidly evolving field. While our study might be relevant and current at the time of writing, the situation could possibly be different in just a few years. Even though the development of the field is rapid, our participants claim that there will be a human operator included in the loop for the foreseeable future. As the future is and will always be uncertain, these limitations should be considered while interpreting the findings of this study.

## 5.4 Suggestions for Further Research

Drawing on the findings and implications of this study, several areas for future research emerge as interesting to enhance our understanding of the interplay between AI and human decision-making. This study has clearly demonstrated that despite AI's increasing influence in organizations, human intervention remains a constant and important component in the decision-making process. This introduces its own set of complexities, one of which includes potential irrational behavior and decision-making by human actors. As such, one possible interesting direction of study would be to explore the emotional responses of end-users towards AI more in depth, as this can significantly affect the acceptance, trust, and ultimately, the value derived from such tools. For instance, negative emotions

experienced by an operator towards an AI tool could engender a sense of mistrust, thereby undermining the perceived efficiency of AI in enhancing decision-making processes.

Similarly, it would be interesting to investigate more cases in which AI has augmented a decision-making process that was previously dependent on heuristics. This would further extend and validate our findings about two of our mechanics cases, providing deeper insights into the likes of end users and customers.

As our results suggest, the organization itself and its culture and routines are equally, if not more, important than the AI tool itself. As such, it would be interesting to further explore the underlying factors of this phenomenon. Further investigation could, therefore, focus on how organizations should design a supportive cultural environment that fosters AI innovation, considering the phenomenon from both the change management perspective, and the learning perspective.

Additionally, given that our interviewees repeatedly underscored the inevitability of human involvement in decision-making processes, future research could seek to further unravel the intricate balance between human and AI responsibility. Consider a scenario where for example a human operator consciously disregards the advice of an AI tool, consequently leading to an unfortunate, and potentially dangerous situation. Or the other way around, where a human operator blindly follows the input from an AI tool, which leads to a similar unfortunate and possibly dangerous situation. In such instances, where should the responsibility lie?

Furthermore, as this is a qualitative study with a limited sample size, additional studies, both qualitative studies with different areas of focus, in addition to quantitative studies, are required to assess the generalization of our findings.

We believe that the exploration of these possible research directions would not only enhance our understanding of the interplay between AI and human decision-

making but would also provide additional guidance to organizations on how to effectively navigate the complexities of AI development and implementation.

## 6.0 Conclusion

This thesis has attempted to address some of the gaps and unanswered questions in the current body of literature, by exploring the integration and implications of Artificial Intelligence in organizational decision-making processes. Through a series of case studies, this paper has demonstrated the practical application of AI based on predictive analytics, NLP, and classification analytics, across various industries. Although the thesis has highlighted that there is no perfect way to utilize AI to augment decision-making processes, the study provides valuable insights into the intricacies an organization is faced with when attempting to develop and implement AI. Based on our findings, there is no doubt that the perceived advantages can exceed the disadvantages of utilizing AI and does so often. It is worth noting, however, that the development and implementation of AI is a costly and time-consuming activity, and that there is always a significant risk of failing no matter how much research is conducted beforehand. Despite the proposed advantages of the technology and the excessive AI hype presented by the media, we want to stress that AI is not magic and increasing organizational knowledge and managing expectations are paramount to companies that aim to augment their decision-making processes with AI.

Moving forward, our research suggests that understanding the full spectrum of AI's impact on decision-making requires a deep exploration of its integration and implications within an organizational context. As such, we call for further research on the emotional responses by end-users towards AI and the usage of other heuristics, as these responses can significantly affect the acceptance, trust, and ultimately, the value derived from such models. Furthermore, we propose that future studies could focus on how organizations should design supportive cultural environments that foster AI innovation, while considering the phenomenon from both a change management perspective, and a learning perspective.

To conclude this dissertation, we want to highlight that while the journey towards AI-augmented decision-making is complex and filled with unavoidable challenges, it also presents numerous opportunities for organizations willing to navigate this minefield. As this is critical for organizational competitiveness, it is vital to escape “the dungeon” in which AI does not “exist”. Although this can be a frightening process, we would like to round off with the optimistic words of Stephen Hawking:

I am an optimist and I believe that we can create AI for the good of the world. That it can work in harmony with us. We simply need to be aware of the dangers, identify them, employ the best possible practice and management, and prepare for its consequences well in advance. (Kharpal, 2017).

## 7.0 References

- Adams, R. D., Bessant, J., & Phelps, R. G. (2006). Innovation management measurement: A review. *International Journal of Management Reviews*, 8(1), 21–47. <https://doi.org/10.1111/j.1468-2370.2006.00119.x>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Press.
- Anyoha, R. (2020, April 23). *The History of Artificial Intelligence*. Harvard University. <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence>
- Arthur, W. B. (1994). *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press. <https://doi.org/10.3998/mpub.10029>
- AthenaTech. (2019, October 21). *AI, Machine Learning (ML) and Natural Language Processing (NLP)*. AthenaTech LLC. <https://athenatech.tech/f/ai-machine-learning-ml-and-natural-language-processing-nlp>
- Barney, J. B. (1991). Firm Resources and Sustained Competitive Advantage. *Journal of Management*, 17(1), 99–120. <https://doi.org/10.1177/014920639101700108>
- Boal, K. & Meckler, M. (2010). Decisions Errors of the 4th, 5th and 6th Kind. In Bolukbasi, T., Chang, K., Zou, J., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Neural Information Processing Systems*, 29, 4356–4364. <https://arxiv.org/pdf/1607.06520>
- Borges, A. F., Laurindo, F. J., Spínola, M. M., Gonçalves, R. F., & Mattos, C. A. (2021). The strategic use of artificial intelligence in the digital era: Systematic literature review and future research directions. *International Journal of Information Management*, 57, 102225. <https://doi.org/10.1016/j.ijinfomgt.2020.102225>
- Bourgeois, L. J., & Eisenhardt, K. M. (1988). Strategic Decision Processes in High Velocity Environments: Four Cases in the Microcomputer Industry. *Management Science*, 34(7), 816–835. <https://doi.org/10.1287/mnsc.34.7.816>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Brown, T., Mann, B. F., Ryder, N. C., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J. C., Winter, C. & Amodei, D. (2020). Language Models are Few-Shot Learners. *Neural Information Processing Systems*, 33, 1877–1901.

- <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>
- Bryman, A., & Bell, E. (2015). *Business Research Methods*. Oxford University Press, USA.
- Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.
- Brynjolfsson, E. & McAfee, A. (2017). *The Business of Artificial Intelligence*. Harvard Business Review. <https://hbr.org/2017/07/the-business-of-artificial-intelligence>
- Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlstrom, P., Henke, N., & Trench, M. (2017). *Artificial intelligence: the next digital frontier?* McKinsey Global Institute. <https://apo.org.au/node/210501>
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1). <https://doi.org/10.1177/2053951715622512>
- Cambridge Dictionary. (2023a). *Tacit knowledge*. <https://dictionary.cambridge.org/dictionary/english/tacit-knowledge>
- Cambridge Dictionary. (2023b). *Explicit knowledge*. <https://dictionary.cambridge.org/dictionary/english/explicit-knowledge>
- Cambridge Dictionary. (n.d.). *Sub-optimization*. <https://dictionary.cambridge.org/dictionary/english/sub-optimization>
- Cannavacciuolo, L., Iandoli, L., Ponsiglione, C., & Zollo, G. (2015). Knowledge elicitation and mapping in the design of a decision support system for the evaluation of suppliers' competencies. *VINE*, 45(4), 530–550. <https://doi.org/10.1108/vine-01-2015-0011>
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nature*, 538(7623), 20–23. <https://doi.org/10.1038/538020a>
- Cath, C. (2018). Governing artificial intelligence: ethical, legal and technical opportunities and challenges. *Philosophical Transactions of the Royal Society A*, 376(2133). <https://doi.org/10.1098/rsta.2018.0080>
- Chandler, A. D. (1962). *Strategy and Structure: Chapters in the History of the American Industrial Enterprise*. <http://ci.nii.ac.jp/ncid/BA77058084>
- Charmaz, K. (2014). *Constructing Grounded Theory* (2nd ed.). Sage Publications.
- Colson, E. (2019). *What AI-Driven Decision Making Looks Like*. <https://hbr.org/2019/07/what-ai-driven-decision-making-looks-like>
- Cremer, D. D. & Kasparov, G. (2021). *AI Should Augment Human Intelligence, Not Replace It*. Harvard Business Review. <https://hbr.org/2021/03/ai-should-augment-human-intelligence-not-replace-it>
- Creswell, J. W. (2014). *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches* (4th ed.). SAGE Publications.



- Creswell, J. W., & Poth, C. (2018). *Qualitative Inquiry and Research Design: Choosing Among Five Approaches* (4th ed.). SAGE Publications.
- Crisan, E. L., & Stanca, L. (2021). The digital transformation of management consulting companies: a qualitative comparative analysis of Romanian industry. *Information Systems and E-business Management*, 19(4), 1143–1173. <https://doi.org/10.1007/s10257-021-00536-1>
- Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship. <https://ssrn.com/abstract=1496208>
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97-118). Oxford University Press.
- Dartnall, T. (1994). Artificial Intelligence and Creativity: An Interdisciplinary Approach. In Kluwer Academic Publishers eBooks. <http://ci.nii.ac.jp/ncid/BA23617659>
- Davenport, T. H., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal*, 6(2), 94–98. <https://doi.org/10.7861/futurehosp.6-2-94>
- Davenport, T. H., & Kirby, J. (2016). Just How Smart Are Smart Machines. *MIT Sloan Management Review*, 57(3), 21–25. <https://dialnet.unirioja.es/servlet/articulo?codigo=5724520>
- Davenport, T. H., & Ronanki, R. (2018). Artificial Intelligence for the Real World. *Harvard Business School*. <https://www.bizjournals.com/boston/news/2018/01/09/hbr-artificial-intelligence-for-the-real-world.html>
- Dean, J. W. & Sharfman, M. P. (1993). Procedural rationality in the strategic decision-making process. *Journal of Management Studies*, 30(4), 587–610. <https://doi.org/10.1111/j.1467-6486.1993.tb00317.x>
- Deci, E. L. & Ryan, R. M. (1985). *Intrinsic Motivation and Self-Determination in Human Behavior*. New York: Plenum.
- Deloitte. (2022). Fueling the AI transformation: Four key actions powering the widespread value from AI, right now. In *Deloitte.com* (5th Edition Report). <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-ai-institute-state-of-ai-fifth-edition.pdf>
- Delua, J. (2021). Supervised vs. Unsupervised Learning: What's the Difference? *IBM*. <https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning>
- Denscombe, M. (2014). *The Good Research Guide: For Small-Scale Social Research Projects* (5th ed.). McGraw-Hill Education (UK).

- Denzin, N. K. (2012). Triangulation 2.0. *Journal of Mixed Methods Research*, 6(2), 80–88. <https://doi.org/10.1177/1558689812437186>
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ACL Anthology*, N19-1423. <https://aclanthology.org/N19-1423.pdf>
- Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition*, 35(7), 1801–1813. <https://doi.org/10.3758/BF03193511>
- Dierkes, M., Antal, A. B., Child, J., & Nonaka, I. (2003). Handbook of Organizational Learning and Knowledge. In Oxford University Press eBooks. <https://ci.nii.ac.jp/ncid/BA61651425>
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J. E., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P. W., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B. & Williams, M. (2021). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- Eisenhardt, K. M. (1989). Building Theories from Case Study Research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.5465/amr.1989.4308385>
- Eisenhardt, K. M., & Sull, D. N. (2001). Strategy as simple rules. *PubMed*, 79(1), 106–116, 176. <https://pubmed.ncbi.nlm.nih.gov/11189455>
- Emanuel, E. J., & Wachter, R. M. (2019). Artificial Intelligence in Health Care. *JAMA*, 321(23), 2281. <https://doi.org/10.1001/jama.2019.4914>
- Engelbart, D. C. (1962). *Augmenting Human Intellect: A Conceptual Framework*. Stanford Research Institute.
- Eschenbach, V., & Warren, J. D. (2021). Transparency and the Black Box Problem: Why We Do Not Trust AI. *Philosophy & Technology*, 34(4), 1607–1622. <https://doi.org/10.1007/s13347-021-00477-0>
- European Parliament. (2023, 14. June). *EU AI Act: first regulation on artificial intelligence*. <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence?fbclid=IwAR1BacT9DN3GuSIelKshKX-oJXEtEtjsjsTa4UCwFv-112Q-TqkTA8H-ZZ4>
- Evans, J. S. (2008). Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition. *Annual Review of Psychology*, 59, 255-278. <https://doi.org/10.1146/annurev.psych.59.103006.093629>
- Ferrer, X., Nuenen, T. v., Such, J. M., Coté, M. & Criado, N. (2021). Bias and Discrimination in AI: A Cross-Disciplinary Perspective. *IEEE Technology and Society Magazine*, 40(2), 72-80. [10.1109/MTS.2021.3056293](https://doi.org/10.1109/MTS.2021.3056293)

- Fieseler C., Bucher E. (2022, 1. November). *Machine Learning (Lecture 5): AI, Algorithms & Society* [PowerPoint Presentation]. Itslearning. <https://bi.itslearning.com/>
- Financial Express. (2023, April 20). Should AI research be paused? *Financial Express*. <https://www.financialexpress.com/opinion/should-ai-research-be-paused/3054551/>
- Fortune Business Insight. (2022). Industry Reports: Artificial Intelligence Market. <https://www.fortunebusinessinsights.com/industry-reports/artificial-intelligence-market-100114>
- Future of Life Institute. (2023, June 29). *Pause Giant AI Experiments: An Open Letter - Future of Life Institute*. <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>
- Galford, R. M., & Drapeau, A. (2003). The enemies of trust. *PubMed*, 81(2), 88–95, 126. <https://pubmed.ncbi.nlm.nih.gov/12577656>
- Gartner. (2018, February 13). *Gartner Says Nearly Half of CIOs Are Planning to Deploy Artificial Intelligence*. <https://www.gartner.com/en/newsroom/press-releases/2018-02-13-gartner-says-nearly-half-of-cios-are-planning-to-deploy-artificial-intelligence>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic Decision Making. *Annual Review of Psychology*, 62(1), 451–482. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Goertzel, B., & Wang, P. (2007). *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms: Proceedings of the AGI Workshop 2006*. IOS Press.
- Grønmo, S. (2004). *Samfunnsvitenskapelige metoder* (1st ed.). Fagbokforlaget.
- Haefner, N., Wincent, J., Parida, V., & Gassmann, O. (2021). Artificial intelligence and innovation management: A review, framework, and research agenda☆. *Technological Forecasting and Social Change*, 162. <https://doi.org/10.1016/j.techfore.2020.120392>
- Haenlein, M., & Kaplan, A. (2019). A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*, 61(4), 5–14. <https://doi.org/10.1177/0008125619864925>
- Hagendorff, T. (2020). The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds and Machines*, 30(1), 99–120. <https://doi.org/10.1007/s11023-020-09517-8>
- Hannan, M. T., & Freeman, J. M. (1984). Structural Inertia and Organizational Change. *American Sociological Review*, 49(2), 149-164. <https://doi.org/10.2307/2095567>
- Harlow, H. D. (2018). Developing a knowledge management strategy for data analytics and intellectual capital. *Meditari Accountancy Research*, 26(3), 400–419. <https://doi.org/10.1108/medar-09-2017-0217>

- Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: matching rules and environments. *Psychological Review*, 114(3), 733–758. <https://doi.org/10.1037/0033-295X.114.3.733>
- Hosmer, L. T. (1995). Trust: The Connecting Link between Organizational Theory and Philosophical Ethics. *Academy of Management*, 20(2), 379–403. <https://www.jstor.org/stable/258851>
- IBM Cloud Education. (2020). *Deep Learning*. <https://www.ibm.com/topics/deep-learning>
- Inbar, Y., Botti, S., & Hanks, K. (2011). Decision speed and choice regret: When haste feels like waste. *Journal of Experimental Social Psychology*, 47(3), 533–540. <https://doi.org/10.1016/j.jesp.2011.01.011>
- Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6), 995–1006. <https://doi.org/10.1037/0022-3514.79.6.995>
- Jacobsen, D.I. & Thorsvik, J. (2019). *Hvordan organisasjoner fungerer* (5. utg.). Fagbokforlaget.
- Jacovi, A., Marasović, A., Miller, T., & Goldberg, Y. (2021). Formalizing Trust in Artificial Intelligence. <https://doi.org/10.1145/3442188.3445923>
- Janiesch, C., Zschech, P. & Heinrich, K. (2021). Machine learning and deep learning. *Electronic Markets*, 31, 685-695. <https://doi.org/10.1007/s12525-021-00475-2>
- Jawad, U. (2023). Ireland data regulator: We should not blindly ban AI-powered chatbots. *Neowin*. <https://www.neowin.net/news/ireland-data-regulator-we-should-not-blindly-ban-ai-powered-chatbots/>
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Madotto, A., & Fung, P. (2023). Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*, 55(12), 1–38. <https://doi.org/10.1145/3571730>
- Jobin, A., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Johannessen, A., Christoffersen, L., & Tufte, P. A. (2020). *Forskningsmetode for økonomisk- administrative fag* (4th ed.). Abstrakt Forlag.
- Jordan, M. I., & Mitchell, T. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Jurafsky, D., & Martin, J. (2019). Information Extraction. In *Speech and Language Processing* (2nd ed.). Prentice Hall. [https://web.stanford.edu/~jurafsky/slp3/old\\_oct19/17.pdf](https://web.stanford.edu/~jurafsky/slp3/old_oct19/17.pdf)
- Kahn, J. (2023, March 29). Elon Musk and Apple cofounder Steve Wozniak among over 1,100 who sign open letter calling for 6-month ban on creating powerful A.I. *Fortune*. <https://fortune.com/2023/03/29/elon-musk-apple->

[steve-wozniak-over-1100-sign-open-letter-6-month-ban-creating-powerful-ai/](#)

- Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, 39(4), 341–350. <https://doi.org/10.1037/0003-066X.39.4.341>
- Kahneman, D. (2011). Thinking, Fast and Slow. Farrar, Straus and Giroux. <https://psycnet.apa.org/record/2011-26535-000>
- Kaplan, A. M., & Haenlein, M. (2019). Siri, Siri, in my hand: Who’s the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25. <https://doi.org/10.1016/j.bushor.2018.08.004>
- Kelleher, J. C., Namee, B. M. & D’Arcy, A. (2020). *Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies* (2nd ed.). MIT Press. [https://books.google.no/books?hl=no&lr=&id=UM\\_tDwAAQBAJ&oi=fnd&pg=PR15&dq=predictive+analytics+machine+learning+introduction&ots=aFOWZqfDi1&sig=fJJWuBN0NICCDWGW752LGC3UAFY&redir\\_esc=y#v=onepage&q&f=false](https://books.google.no/books?hl=no&lr=&id=UM_tDwAAQBAJ&oi=fnd&pg=PR15&dq=predictive+analytics+machine+learning+introduction&ots=aFOWZqfDi1&sig=fJJWuBN0NICCDWGW752LGC3UAFY&redir_esc=y#v=onepage&q&f=false)
- Kharpal, A. (2017). *Stephen Hawking says A.I. could be ‘worst event in the history of our civilization’*. CNBC. <https://www.cnbc.com/2017/11/06/stephen-hawking-ai-could-be-worst-event-in-civilization.html>
- Khurana, D., Koli, A. C., Khatter, K., & Singh, S. (2022). Natural language processing: state of the art, current trends and challenges. *Multimedia Tools and Applications*. <https://doi.org/10.1007/s11042-022-13428-4>
- Kolbjørnsrud, V., Amico, R., & Thomas, R. J. (2016). How Artificial Intelligence Will Redefine Management. Harvard Business Review. [https://www.researchgate.net/publication/311775855\\_How\\_Artificial\\_Intelligence\\_Will\\_Redefine\\_Management](https://www.researchgate.net/publication/311775855_How_Artificial_Intelligence_Will_Redefine_Management)
- Kolbjørnsrud, V., Amico, R., & Thomas, R. (2017). Partnering with AI: how organizations can win over skeptical managers. *Strategy & Leadership*, 45(1), 37–43. <https://doi.org/10.1108/sl-12-2016-0085>
- Kotter, J. P. (1995, May-June). Leading Change: Why Transformation Efforts Fail. *Harvard Business Review*. <https://hbr.org/1995/05/leading-change-why-transformation-efforts-fail-2>
- Kumar, V., & Garg, M. (2018). Predictive Analytics: A Review of Trends and Techniques. *International Journal of Computer Applications*, 182(1), 31–37. <https://doi.org/10.5120/ijca2018917434>
- Lane, M., Williams, M., & Broecke, S. (2023). The impact of AI on the workplace: Main findings from the OECD AI surveys of employers and workers. In *OECD* (No. 288). Organisation for Economic Co-operation and Development (OECD). <https://www.oecd-ilibrary.org/social-issues->

- migration-health/the-impact-of-ai-on-the-workplace-main-findings-from-the-oecd-ai-surveys-of-employers-and-workers\_ea0a0fe1-en
- Langley, A. (1995, April 15). Between “Paralysis by Analysis” and “Extinction by Instinct.” *MIT Sloan Management Review*.  
<https://sloanreview.mit.edu/article/between-paralysis-by-analysis-and-extinction-by-instinct/>
- Langley, A., Mintzberg, H., Pitcher, P., Posada, E., & Saint-Macary, J. (1995). Opening up Decision Making: The View from the Black Stool. *Organization Science*, 6(3), 260–279.  
<http://www.jstor.org/stable/2635251>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Legg, S., & Hutter, M. (2007). Universal Intelligence: A Definition of Machine Intelligence. *Minds and Machines*, 17(4), 391–444.  
<https://doi.org/10.1007/s11023-007-9079-x>
- Leonard-Barton, D. & Kraus, W. A. (1985). Implementing New Technology. *Harvard Business School*. <https://hbr.org/1985/11/implementing-new-technology>
- Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. *International Journal of Information Management*, 50, 57–70.  
<https://doi.org/10.1016/j.ijinfomgt.2019.04.003>
- Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and Decision Making. *Annual review of psychology*, 66, 799–823.  
<https://doi.org/10.1146/annurev-psych-010213-115043>
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic Inquiry*. SAGE.  
<https://us.sagepub.com/en-us/nam/naturalistic-inquiry/book842>
- Littman, M. L., Ajunawa, I., Berger, G., Boutilier, C., Currie, M., Doshi-Velez, F., Hadfield, G., Horowitz, M. C., Isbell, C., Kitano, H., Levy, K., Lyons, T., Mitchell, M., Sloman, S., Shah, J., Vallor, S., & Walsh, T. (2021, September). Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report. Stanford University. <https://ai100.stanford.edu/gathering-strength-gathering-storms-one-hundred-year-study-artificial-intelligence-ai100-2021-study>
- Lones, M. A. (2021). How to avoid machine learning pitfalls: a guide for academic researchers. In Arxiv. Cornell University.  
<https://arxiv.org/abs/2108.02497>
- March, J. G. (1978). Bounded Rationality, Ambiguity, and the Engineering of Choice. *The Bell Journal of Economics*, 9(2), 587-608.  
<https://doi.org/10.2307/3003600>
- March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Studies*, 2(1), 71-87.<https://www.jstor.org/stable/2634940>

- March, J. G. (1994). *A Primer on Decision Making: How Decisions Happen*. Free Press, New York.
- Mathew, D., Shukla, V. K., Chaubey, A., & Dutta, S. (2021). Artificial Intelligence: Hope for Future or Hype by Intellectuals? 2021 9th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO). <https://doi.org/10.1109/icrito51393.2021.9596410>
- Maxwell, J. A. (2013). *Qualitative Research Design: An Interactive Approach* (3rd ed.). SAGE.
- McCallum, B. S. (2023, April 1). *ChatGPT banned in Italy over privacy concerns*. BBC News. <https://www.bbc.com/news/technology-65139406>
- McCarthy, J. J., Minsky, M., Rochester, N., & Shannon, C. E. (1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>
- McKinsey & Company. (2022, December 6). *The state of AI in 2022 - and a half decade in review*. McKinsey.com. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review#/>
- Metcalf, L., Askay, D. A., & Rosenberg, L. B. (2019). Keeping Humans in the Loop: Pooling Knowledge through Artificial Swarm Intelligence to Improve Business Decision Making. *California Management Review*, 61(4), 84–109. <https://doi.org/10.1177/0008125619862256>
- Milkman, K. L., Chugh, D., & Bazerman, M. H. (2009). How Can Decision Making Be Improved? *Perspectives on Psychological Science*, 4(4), 379–383. <https://doi.org/10.1111/j.1745-6924.2009.01142.x>
- MIT [Lex Fridman]. (2019). *Deep Learning Basics: Introduction and Overview* [Video]. YouTube. <https://www.youtube.com/watch?v=O5xeyoRL95U>
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507. <https://doi.org/10.1038/s42256-019-0114-4>
- Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). <https://doi.org/10.1177/2053951716679679>
- Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2021). Operationalising AI Ethics: Barriers, Enablers and Next Steps. *AI & Society*, 38(1), 411–423. <https://doi.org/10.1007/s00146-021-01308-8>
- Mollick, E. (2022). ChatGPT Is a Tipping Point for AI. *Harvard Business School*. <https://hbr.org/2022/12/chatgpt-is-a-tipping-point-for-ai>
- Nader, K., Toprac, P., Scott, S. E., & Baker, S. W. (2022). Public understanding of artificial intelligence through entertainment media. *AI & Society*. <https://doi.org/10.1007/s00146-022-01427-w>
- Nadkarni, P. M., Ohno-Machado, L., & Chapman, W. W. (2011). Natural language processing: an introduction. *Journal of the American Medical*

- Informatics Association*, 18(5), 544–551. <https://doi.org/10.1136/amiajnl-2011-000464>
- Nelson, R. E. & Winter, S. G., & (1983). An Evolutionary Theory of Economic Change. *The Economic Journal*, 93(371), 652.  
<https://doi.org/10.2307/2232409>
- Nilsson, N. J. (2010). *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Web Version ed.). Cambridge.  
<https://ai.stanford.edu/~nilsson/QAI/qai.pdf>
- Nishant, R., Kennedy, M. & Corbett, J. (2020). Artificial intelligence for sustainability: Challenges, opportunities, and a research agenda. *International Journal of Information Management*, 53.  
<https://doi.org/10.1016/j.ijinfomgt.2020.102104>
- Nutt, P. C. (2008). Investigating the Success of Decision Making Processes. *Journal of Management Studies*, 45(2), 425-455.  
<https://doi.org/10.1111/j.1467-6486.2007.00756.x>
- Nutt, P. C. (2011). Making Decision-Making Research Matter: Some Issues and Remedies. *Management Research Review*, 34(1), 5-16.  
<https://doi.org/10.1108/01409171111096441>.
- OpenAI. (2023a). *Introducing ChatGPT*. <https://openai.com/blog/chatgpt>
- OpenAI. (2023b). GPT-4 Technical Report. In *Cdn.openai.com*.  
<https://cdn.openai.com/papers/gpt-4.pdf>
- Papadakis, V. M., Lioukas, S., & Chambers, D. (1998). Strategic decision-making processes: the role of management and context. *Strategic Management Journal*, 19(2), 115–147. [https://doi.org/10.1002/\(sici\)1097-0266\(199802\)19:2](https://doi.org/10.1002/(sici)1097-0266(199802)19:2)
- Patton, M. Q. (2015). *Qualitative Research & Evaluation Methods: Integrating Theory and Practice* (4th ed.). SAGE Publications.
- Payne, J. W., Bettman, J. R. & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
- Perifanis, N.-A., & Kitsios, F. (2023). Investigating the Influence of Artificial Intelligence on Business Value in the Digital Era of Strategy: A Literature Review. *Information*, 14(2), 85. <https://doi.org/10.3390/info14020085>
- Pfiffner, J. M. (1960). Administrative Rationality. *Public Administration Review*, 20(3), 125-132. <https://doi.org/10.2307/973965>
- Phillips-Wren, G., & Adya, M. (2020). Decision making under stress: the role of information overload, time pressure, complexity, and uncertainty. *Journal of Decision Systems*, 29(1), 213–225.  
<https://doi.org/10.1080/12460125.2020.1768680>
- Porter, M. E. (1985). *Competitive advantage: Creating and sustaining superior performance*. Free Press, New York Press.  
[http://www.dangoldstein.com/papers/Czerlinski\\_How\\_Good\\_Are\\_Simple\\_Heuristics\\_Chapter\\_1999.pdf](http://www.dangoldstein.com/papers/Czerlinski_How_Good_Are_Simple_Heuristics_Chapter_1999.pdf)



- Qu, S. Q., & Dumay, J. (2011). The qualitative research interview. *Qualitative Research in Accounting & Management*, 8(3), 238–264. <https://doi.org/10.1108/11766091111162070>
- Ridder, H. (2017). The theory contribution of case study research designs. *Business Research*, 10(2), 281–305. <https://doi.org/10.1007/s40685-017-0045-z>
- Robertson, A. (2023, April 28). ChatGPT returns to Italy after ban. *The Verge*. <https://www.theverge.com/2023/4/28/23702883/chatgpt-italy-ban-lifted-gdp-data-protection-age-verification>
- Roser, M. (2023, February 7). *AI timelines: What do experts in artificial intelligence expect for the future?* Our World in Data. <https://ourworldindata.org/ai-timelines>
- Russell, S. (2023, March 27-29). *The shifting landscape of AI* [Conference presentation]. International Conference on AI in Work, Innovation, Productivity and Skills, Online.
- Russell, S., & Norvig, P. (2009). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson Education. [https://people.engr.tamu.edu/guni/csce421/files/AI\\_Russell\\_Norvig.pdf](https://people.engr.tamu.edu/guni/csce421/files/AI_Russell_Norvig.pdf)
- Sarker, I. H. (2021a). Machine Learning: Algorithms, Real-World Applications and Research Directions. *SN Computer Science*, 2(3). <https://doi.org/10.1007/s42979-021-00592-x>
- Sarker, I. H. (2021b). Data Science and Analytics: An Overview from Data-Driven Smart Computing, Decision-Making and Applications Perspective. *SN Computer Science*, 2(5). <https://doi.org/10.1007/s42979-021-00765-8>
- Scarpetta, S. (2023, March 27-29). *The shifting landscape of AI* [Conference presentation]. International Conference on AI in Work, Innovation, Productivity and Skills, Online.
- Schilling, M. A., & Phelps, C. (2007). Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation. *Management Science*, 53(7), 1113–1126. <https://doi.org/10.1287/mnsc.1060.0624>
- Schoemaker, J. H., & Tetlock, P. E. (2017, March 13). Building a More Intelligent Enterprise. *MIT Sloan Management Review*. Retrieved January 3, 2023, from <https://sloanreview.mit.edu/article/building-a-more-intelligent-enterprise/>
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, 134(2), 207–222. <https://doi.org/10.1037/0033-2909.134.2.207>
- Shepherd, N. G., & Rudd, J. M. (2014). The Influence of Context on the Strategic Decision-Making Process: A Review of the Literature. *International Journal of Management Reviews*, 16(3), 340–364. <https://doi.org/10.1111/ijmr.12023>

- Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. <https://mitpress.mit.edu/9780262515474/telerobotics-automation-and-human-supervisory-control/>
- Shrestha, Y. R., Ben-Menahem, S. M., & Von Krogh, G. (2019). Organizational Decision-Making Structures in the Age of Artificial Intelligence. *California Management Review*, 61(4), 66–83. <https://doi.org/10.1177/0008125619862257>
- Siegel, E. (2023, 2. June). The AI Hype Cycle Is Distracting Companies. *Harvard Business School*. <https://hbr.org/2023/06/the-ai-hype-cycle-is-distracting-companies>
- Simon, H. A. (1947). *Administrative behaviour: A Study of Decision-Making Processes in Administrative Organization*. New York: Macmillan.
- Simon, H. A. (1957). *Models of man, social and rational: mathematical essays on rational human behavior in a social setting*. Published in 1957 in New York by Wiley.
- Simon, H. A. (1960). *The New Science of Management Decision* (1st ed.). Harper & Borthers Publishers. <https://babel.hathitrust.org/cgi/pt?id=uc1.b3376401&view=1up&seq=7&q1=design>
- Simon, H. A. (1964). On the Concept of Organizational Goal. *Administrative Science Quarterly*, 9(1), 1-22. <https://doi.org/10.2307/2391519>
- Simon, H. A. (1966). *The New Science of Management Decision*. New York: Harper.
- Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, 74(1), 29-39. <https://doi.org/10.1037/h0024127>
- Simon, H. A. (1973). Applying information technology to organization design. *Public Administration Review*, 33(3), 268-278. <https://www.jstor.org/stable/974804>
- Simon, H. A. (1979). Rational decision making in business organizations. *The American Economic View*, 69(4), 493-513. <https://www.jstor.org/stable/1808698>
- Slovic, P., Finucane, M. L., Peters, E. & MacGregor, D. G. (2006). The affect heuristic. *European Journal of Operational Research*, 177(3), 1333-1352. <https://doi.org/10.1016/j.ejor.2005.04.006>
- Smith, J. A., Flowers, P., & Larkin, M. (2009). *Interpretative Phenomenological Analysis: Theory, Method and Research*. SAGE.
- Snow, C. C., & Fjeldstad, O. D. (2023). Organizing Intelligent Digital Actors. *Designing Adaptive Organizations*.
- Sonenshein, S. (2007). The role of construction, intuition, and justification in responding to ethical issues at work: The sensemaking-intuition model. *Academy of Management Review*, 32(4), 1022–1040. <https://doi.org/10.5465/amr.2007.26585677>

- Stephoe-Warren, G., Howat, D., & Hume, I. (2011). Strategic thinking and decision making: literature review. *Journal of Strategy and Management*, 4(3), 238–250. <https://doi.org/10.1108/17554251111152261>
- Stobierski, J. (2019, August 26). The Advantages of Data-Driven Decision Making. *Harvard Business School Online*. Retrieved January 5, 2023, from <https://online.hbs.edu/blog/post/data-driven-decision-making>
- Suri, H. (2011). Purposeful Sampling in Qualitative Research Synthesis. *Qualitative Research Journal*, 11(2), 63–75. <https://doi.org/10.3316/qrj1102063>
- Sydow, J., Schreyögg, G. & Koch, J. (2009). Organizational Path Dependence: Opening the Black Box. *Academy of Management Review*, 34(4), 689-709. <https://doi.org/10.5465/amr.34.4.zok689>
- Tam, G. K. L., Kothari, V., & Chen, M. (2017). An Analysis of Machine- and Human-Analytics in Classification. *IEEE Transactions on Visualization and Computer Graphics*, 23(1), 71–80. <https://doi.org/10.1109/tvcg.2016.2598829>
- The National Committee for Research Ethics in the Social Sciences and the Humanities. (2022, May 26). *Guidelines for Research Ethics in the Social Sciences and the Humanities*. Forskningsetikk. <https://www.forskningsetikk.no/en/about-us/our-committees-and-commission/nesh/guidelines-nesh/guidelines-for-research-ethics-in-the-social-sciences-humanities-law-and-theology/>
- Trivedi, U. B., Bhatt, M. B., & Srivastava, P. (2021). Prevent Overfitting Problem in Machine Learning: A Case Focus on Linear Regression and Logistics Regression. In *Advances in science, technology & innovation* (pp. 345–349). Springer International Publishing. [https://doi.org/10.1007/978-3-030-66218-9\\_40](https://doi.org/10.1007/978-3-030-66218-9_40)
- Trocin, C., Hovland, I. V., Mikalef, P. & Dremel, C. (2021). How Artificial Intelligence affords digital innovation: A cross-case analysis of Scandinavian companies. *Technological Forecasting and Social Change*, 173, 1-12. <https://doi.org/10.1016/j.techfore.2021.121081>
- Turing, A. M. (1950). *Computing Machinery and Intelligence*. *Mind*, LIX(236), 433–460. <https://doi.org/10.1093/mind/lix.236.433>
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- United Nations. (2022). *Goal 13: Take urgent action to combat climate change and its impacts*. <https://sdgs.un.org/goals/goal13>
- University of Southern California. (n.d.). *Research Guides: Organizing Your Social Sciences Research Paper*. USC Libraries. <https://libguides.usc.edu/writingguide/introduction/researchproblem>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is All you Need. In arXiv

- (Cornell University) (Vol. 30, pp. 5998–6008). Cornell University.  
<https://arxiv.org/pdf/1706.03762v5>
- Wang, P. (1995). *On the Working Definition of Intelligence*. Center for Research on Concepts and Cognition.  
[https://www.researchgate.net/publication/2339604\\_On\\_the\\_Working\\_Definition\\_of\\_Intelligence](https://www.researchgate.net/publication/2339604_On_the_Working_Definition_of_Intelligence)
- Weick, K. E. (2009). *Making Sense of the Organization. The Impermanent Organization*. Chichester: John Wiley.
- Wilks, Y. (2019). *Artificial Intelligence: Modern Magic or Future*. HotScience.  
[https://books.google.no/books?hl=no&lr=&id=966WDwAAQBAJ&oi=fnd&pg=PT5&dq=artificial+intelligence+dangers&ots=EiQ3FEZD5&sig=WhJwUdvXEtpUFERp6MZTIc0UcM&redir\\_esc=y#v=onepage&q=artificial%20intelligence%20dang&f=false](https://books.google.no/books?hl=no&lr=&id=966WDwAAQBAJ&oi=fnd&pg=PT5&dq=artificial+intelligence+dangers&ots=EiQ3FEZD5&sig=WhJwUdvXEtpUFERp6MZTIc0UcM&redir_esc=y#v=onepage&q=artificial%20intelligence%20dang&f=false)
- Wolfe, C. R. (2022). *Language Models: GPT and GPT-2*. Towards Data Science.  
<https://towardsdatascience.com/language-models-gpt-and-gpt-2-8bdb9867c50a>
- Xiao, D., & Su, J. (2022). Role of Technological Innovation in Achieving Social and Environmental Sustainability: Mediating Roles of Organizational Innovation and Digital Entrepreneurship. *Frontiers in Public Health, 10*.  
<https://doi.org/10.3389/fpubh.2022.850172>
- Yin, R. K. (2009). *Case Study Research: Design and Methods* (4th ed.). SAGE.  
[https://books.google.no/books?hl=no&lr=&id=FzawIAdilHkC&oi=fnd&pg=PR1&dq=Case+Study+Research:+Design+and+Methods&ots=l\\_5T3diWYq&sig=hJbPkRgcOjfC6O5IRzYsCk64n-8&redir\\_esc=y#v=onepage&q=Case%20Study%20Research%3A%20Design%20and%20Methods&f=false](https://books.google.no/books?hl=no&lr=&id=FzawIAdilHkC&oi=fnd&pg=PR1&dq=Case+Study+Research:+Design+and+Methods&ots=l_5T3diWYq&sig=hJbPkRgcOjfC6O5IRzYsCk64n-8&redir_esc=y#v=onepage&q=Case%20Study%20Research%3A%20Design%20and%20Methods&f=false)
- Yukl, G. & Lepsinger, R. (2006). Leading Change: Adapting and Innovating in an Uncertain World. *Leadership in Action, 26*(2), 3-7.  
<https://doi.org/10.1002/lia.1154>

## 8.0 Appendix

### Semi-Structured Interview Guide

#### **Introduction:**

Thanks for meeting us, we really appreciate it. Just before we start: We would like to record the conversation to transcribe and categorize the data after the interview. This will only be used in connection with our master's thesis, and all audio recordings and associated transcription documents will be stored securely and deleted by the end of the project on July 3, 2023. Is it okay that we record and transcribe the conversation?

-> If okay, start recording.

We would also like to remind you that participation is voluntary and that you can withdraw from the project at any time if you want. Furthermore, we anticipate that the interview will take up to one hour.

Do you have any questions before we begin?

**Explanation of:** Who we are, what are we looking into and what do we aim to get out of this interview.

#### **Questions:**

##### Introductory Questions

1. Who are you, and what is your role? How long have you been working in the company?
2. Briefly describe the company/department you work for.

##### Topic – AI

3. What does artificial intelligence mean to you? What is your definition?
4. How do you perceive the attitude within the company regarding the implementation of AI?
5. What is your attitude?
6. Has AI been adopted in other parts of the company? Are there different types of AI?
7. How many people are working with AI in the organization?

Topic – AI in Decision-Making Processes

8. Can you describe a case where you have used AI?
9. How did the project start? Who started it? What was the motivation behind it?
10. How long have you been working on the case? What is the status of the project?
11. Can you describe the decision-making process in the case as it was before the use of AI?
  - How did you experience the decision-making process here?
12. Can you describe the decision-making process in the case as it is now (after implementation)?
  - How did you experience the decision-making process here?
13. What are the advantages of using AI in the decision-making process?
  - Higher strategic goal achievement?
  - How does the implementation of AI align with the company/group strategy?
  - Increased task efficiency? Speed?
  - Automation?
  - Cost savings?
  - More rational decision-making? (Explain rationality if not known)
  - Sustainability?
  - Does the decision-making process improve? Better decisions?
14. What are the disadvantages/limitations of using AI in the decision-making process?
  - Explanation problems (transparency and trust)
  - Biases
15. How are incorrect answers addressed during training?
16. What obstacles have you encountered in implementing AI in the decision-making process?
17. How do you ensure that the use of AI does not negatively affect privacy and data security?

18. Have you received feedback from customers or other stakeholders regarding the use of AI?

19. Is the use of AI publicly visible?

20. How do you envision this process in the future with more advanced technology and models?

21. Do you have any tips/suggestions for other organizations considering implementing AI in their processes?

Concluding Questions

22. Is there anything we haven't asked about that you think is relevant to mention?

23. Is there anything else you would like to know about our thesis?

24. Do you know anyone else we should talk to? Someone directly affected by the implementation of AI, i.e., Managers, colleagues, customers, or others?