



Value Creation Opportunities of Generative AI – A Case Study

Alexander Sake

Technical University of Munich

Abstract

The transformative potential of Generative AI promises novel capabilities within business environments. This study examines the value creation potential of Generative AI within a large multinational corporation. A single case study approach at Siemens was employed, combining extensive observations, interviews, and the application of existing AI frameworks. Findings reveal diverse use cases demonstrating value creation potential, particularly through smart assistants and lighthouse projects. This thesis proposes a novel framework for Generative AI adoption, emphasizing the distinctive exploration phase made possible by the technology's accessibility to non-technical domain experts, while also outlining essential scaling strategies. This study offers valuable insights into a company's approach to Generative AI, provides practical implications, and expands ongoing research on AI-driven value creation.

Keywords: artificial intelligence; exploration; generative AI; scaling; technology adoption; use cases; value creation

1. Introduction

The release of ChatGPT in November 2022 marked a significant moment where the public could, for the first time, directly engage with the latest advances in artificial intelligence (AI). Since its launch, the tool's impact and capabilities have been a topic of diverse discussions among journalists, scientists, managers, and governments, oscillating between praise and cautionary notes. This thesis intends to shed light on the latest advances of the technology, to describe the value creation opportunities that arise from it and to help companies identifying requisites and capabilities for a successful introduction into their organization.

Many studies have been conducted on the impact of AI on companies and society. Following technological advance-

ments, these studies often predicted new value creation opportunities arising from the successful implementation of AI solutions in organizations. Missing out on AI, results frequently indicated a significant gap in future performance and competitiveness. A study highlighting these aspects, was conducted by Accenture (Reilly et al., 2019). Based on a global survey among 1,500 C-suite executives, they discovered that 84% of leaders believe to achieve their growth ambitions only with the help of AI. At the same time, 76% struggle to scale the technology across their organization.

The rise of Generative AI was accompanied by many studies as well, evaluating real-world use cases and analyzing the impact on organizational processes and the future of work. In particular, when the GPT-4 model was released, questions arose how the tool will transform and potentially replace activities of knowledge workers (Dwivedi et al., 2023, p. 7). A study by BCG among 750 consultants revealed a performance increase of 40% when using GPT-4 for typical consulting activities compared to a control group (Dell'Acqua et al., 2023, p. 17). This was accompanied by an increase in speed of 25% with a positive impact across all skill levels. Focusing on the performance of customer support agents, Brynjolfsson et al. (2023) also examined the impact of Genera-

I would like to express my sincere gratitude to my supervisor, Tim Lämmermann, M.Sc., for his valuable support and guidance throughout my master thesis. I also extend my special thanks to Prof. Dr. Thomas Hutzschenreuter, Chair of Strategic and International Management, for his inspiring teaching and expertise. I am deeply grateful to all the interview partners who contributed to my work with their valuable insights and experiences. A big thank you goes to my managers for allowing me to pursue this project alongside my job. Lastly, I would like to thank my family and friends – your understanding and support during the intense period of the master thesis were invaluable.

tive AI. Based on data from 5,176 agents, they found that AI was able to increase productivity by 14%. In their survey, differences in regard to skill-levels became apparent: While new or low-skilled workers' performance increased by 34%, highly skilled workers barely benefited from the tool (2023, p. 1). While these examples show the potential added value in business settings, other studies focus on the impact on job roles and the global economy. Findings by Goldman Sachs Research suggest that Generative AI could have a "profound" effect on the world economy and society, potentially raising global GDP by 7% over a 10-year period (Goldman Sachs, 2023). A report by Eloundou et al. (2023) focusing on the US labor market see 80% of the workforce impacted by Generative AI by at least 10% while for almost 20% of the workforce, half of the tasks are seen to be exposed to AI (2023, p. 1).

2. Research Objective

Generative AI represents a significant advancement in artificial intelligence, characterized by its ability to produce new content through pattern recognition in existing data (Feuerriegel et al., 2023, p. 1). This is a shift from traditional AI, which primarily dealt with tasks such as identification and classification (LeCun et al., 2015, p. 436). A major development in Generative AI was the introduction of the Transformer architecture in 2017 (Vaswani et al., 2017), which enhanced the AI's capability to comprehend and process extensive information, thereby improving its content generation ability. Further advancements were achieved with the introduction of the Generative Pre-trained Transformer (GPT) model by Radford and Narasimhan (2018), representing a significant step forward in the field of AI-driven text generation. Trained on diverse datasets, the model was capable of producing text that is not only coherent but also contextually relevant.

The primary **Research Question (RQ1)** intends to delve into the advancements introduced by Vaswani et al. (2017) and Radford and Narasimhan (2018). It aims to investigate how Generative AI diverges from previous AI models in terms of learning, information processing, and content creation. Grasping these attributes of Generative AI is key to understanding its transformative impact. To address these considerations, the research question is formulated as follows:

RQ1: How does the architecture and functionality of Generative AI differ from previous AI models?

The outlined studies conducted by Brynjolfsson et al. (2023) and Dell'Acqua et al. (2023) provide initial insights into the potential of Generative AI to enhance workforce productivity. However, further research is necessary to gain a deeper understanding of the technology and its potential benefits, particularly for larger organizations. The disruptive nature of AI has long been a subject of investigation, yet many companies struggle to translate AI's potential into

tangible business value (Shollo et al., 2022, p. 1). To address this gap, it is crucial to analyze the characteristics of Generative AI and explore how it can benefit larger corporations on a broader scale. This exploration is essential for identifying and understanding the opportunities that Generative AI presents for creating value within large corporate environments. By doing so, the thesis aims to contribute to the ongoing discussion on the value creation potential of AI. Expressed in **Research Question 2:**

RQ2: To what extent does Generative AI open up new opportunities for value creation in companies?

Based on the potential value-add, the question arises how companies can organize and drive the implementation of AI in their organization. Previous research has shown that multiple dimensions need to be considered to adopt AI successfully (Uren & Edwards, 2023). However, depending on the focus area of the research, different requirements have been proposed. While some studies focus on technological aspects of AI (Davenport & Ronanki, 2018, p. 52), others highlight the importance of strategic considerations (Brock & Wangenheim, 2019, p. 7). Providing insights from a case study could help to verify the various aspects of these studies, strengthen the outlined models, and add to the ongoing discourse in the field of science. Articulated in **Research Question 3:**

RQ3: What considerations should companies make in order to exploit the value creation potential of Generative AI?

The thesis aims to provide valuable insights for both scholars and professionals through exploring the defined research questions. Firstly, it seeks to evaluate existing models in real-world scenarios, helping to understand their practical use and how earlier research relates to Generative AI. Secondly, it aims to offer practitioners useful insights into potential uses and strategies for scaling Generative AI in their organizations.

In concluding this chapter, it is vital to highlight that a comprehensive interpretation of the term Generative AI should be maintained throughout the thesis. Given the exploratory character of the case study and the novelty of the technology, in the interviews, a clear distinction between different technological aspects was not always made. This holds in particular true for the terms "Generative AI", "ChatGPT" and "Large Language Model (LLM)", which have been frequently used. As this work is about identifying value creation opportunities from a business perspective rather than detailing exact technical mechanisms, this should be considered an acceptable inaccuracy.

3. Conceptual and Theoretical Background

The next chapter delves into the conceptual and theoretical framework essential for the thesis. It begins by highlighting key developmental stages that have led to Generative AI.

Following that, it introduces theoretical concepts that shed light on how AI can enhance value creation and the essential capabilities required by companies to achieve this.

3.1. Evolution of Artificial Intelligence and Machine Learning

The idea of artificial intelligence dates back to the 1950s when the mathematician and computer scientist Alan Turing asked himself how the intelligence of computers could be measured (Turing, 1950). Instead of asking whether a machine is intelligent, he introduced the “Imitation Game” (now known as the Turing Test) to ask whether machines can imitate human responses convincingly (Turing, 1950, p. 433). Turing described machines or digital computers as complex systems capable of a wide range of tasks, similar to a human following instructions (Turing, 1950, p. 436). He suggested that these computers, due to their vast capabilities, could potentially pass as human in his test (Turing, 1950, p. 442). This idea was a significant step in understanding machine intelligence, proposing a practical way to measure it and laying the groundwork for the field of artificial intelligence.

The term “Artificial Intelligence” itself only became a collective term for a variety of different concepts in 1956. McCarthy, Assistant Professor of Mathematics at Dartmouth College, chose the term for a workshop on the topic and is therefore regarded today, alongside Marvin Minsky, Allen Newell, Herbert Simon as one of the “fathers of AI” (Nilsson, 2013, p. 80). McCarthy was among the first who outlined the requirements for a system to evolve human-like intelligence in more detail. In his paper “programs with common sense” (1958), he discussed features which would be required for a machine to evolve intelligence. These included the capability to represent all behaviors, simply express interesting behavioral changes, improve most aspects of behavior, understand partial success in complex problems, and create improvable subroutines (McCarthy, 1958, p. 5).

Building on McCarthy’s foundational work, various definitions of AI have emerged over time. Simmons and Chappell (1988, p. 14) define AI as the behavior of a machine which, if a human behaves in the same way, would be considered intelligent. Luger and Stubblefield (1998, p. 1) define AI as a branch of computer science focused on the automation of intelligent behavior. In contrast, Russell and Norvig (2021, p. 2) follow an rational agent approach, defining AI as the study and construction of intelligent agents. Based on the scientific work, the following definition shall be applied in this thesis: AI is defined as the ability of machines or computers to learn and perform tasks that are typically attributed to human intelligence.

The idea of teaching machines to learn can be traced back to Samuel (1959). Using the game of checkers, he was able to program a computer that was able to play the game and outperform a human player based on a rudimentary set of parameters and rules (Samuel, 1959, p. 535). His novel concept of machines that could improve their performance over time through experience, is often cited as one of the earliest

work of “machine learning” (McCarthy & Feigenbaum, 1990, p. 10).

What does “learning” actually mean when discussing computer programs? In his book “Machine Learning” (1997), T. Mitchell offers the following definition: “A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E ” (Mitchell, 1997, p. 2). He also uses the game of checkers as an example to describe a computer program that improves its ability to win by gaining experience. This example can be applied to more than simple games. More complex rules and sequences of instructions, better known as algorithms, now enable computer programs to learn a wide variety of tasks. Speech recognition, driving autonomous cars and the classification of new astronomical structures are examples where machine learning is used in practice today (Mitchell, 1997, p. 3).

Before delving into the specifics of machine learning, **Figure 1** presents key terms that are beneficial for subsequent exploration of the topic. Machine learning, which can be seen as a subfield of artificial intelligence, encompasses diverse learning paradigms. These will be outlined in the following section, preceding in-depth explorations of deep learning and Generative AI in later chapters.

The first type of learning is called supervised learning (Goodfellow et al., 2016, p. 103). In this method, the machine or algorithm is provided with a dataset whose examples are labeled. When new, unlabeled data is presented, the machine attempts to predict the label based on the patterns and properties of the known data. The result is then compared with the actual label. The training of the algorithm consists of reducing the error between the estimated and the actual labels, enabling the machine to recognize new input and classify it correctly (LeCun et al., 2015, p. 436). The second type is called self- or unsupervised learning. Instead of providing the machine with labeled data, it learns to identify patterns and properties from the data itself. This way, it learns about the probability distribution of the entire dataset (Goodfellow et al., 2016, p. 103). While this type of learning has been explored for a long time (see Hinton and Sejnowski, 1985), it gained popularity with the introduction of Transformer architecture, large datasets, and the availability of computing power (Radford et al., 2019, p. 10). A third type is called reinforcement learning, which introduces a feedback loop into the learning process. During training, the model uses its experiences to improve performance through a reward and penalty system, aiming to increase its cumulative rewards over time (Sutton & Barto, 2018, p. 2).

3.1.1. From Turing to Deep Learning

Over the past few decades, research in the field of AI has made substantial progress. New methodologies, models, and architectures have been developed, vastly surpassing initial conceptions. These advancements enable AI to perform tasks such as writing poems, developing software, and composing music (Feuerriegel et al., 2023, p. 1). A pivotal technique

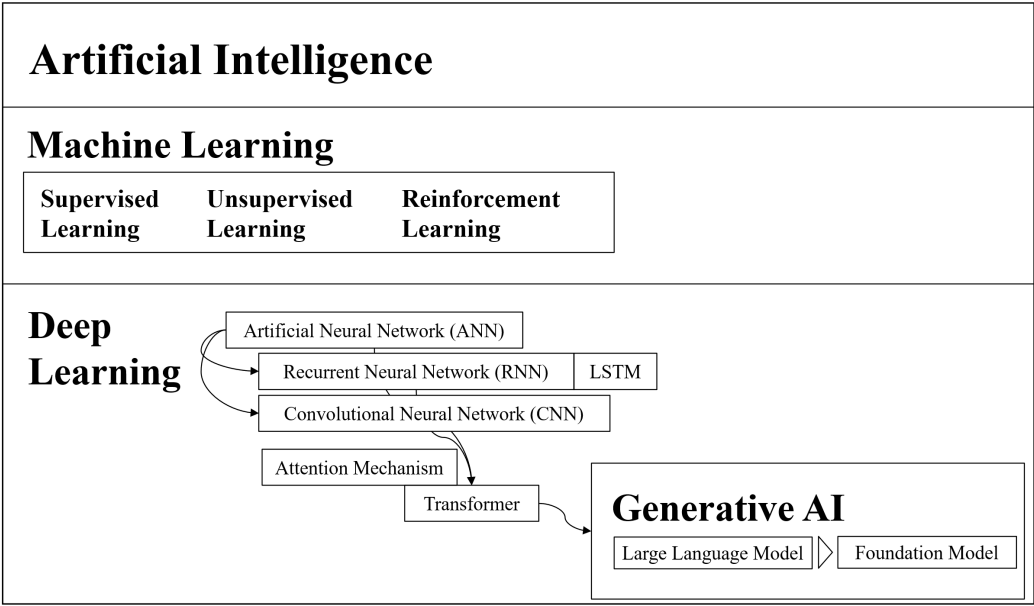


Figure 1: Relationship between AI terms which will be discussed in this thesis.

that has been central to these capabilities is deep learning (Goodfellow et al., 2016, p. 5), which involves decomposing complex relationships into simpler, interconnected concepts. These simpler concepts can then be represented by even more fundamental concepts (Goodfellow et al., 2016, p. 1). In this way, deep learning enables computers to infer complex relationships from simpler ones (Goodfellow et al., 2016, p. 2). The architecture of these models comprises numerous hierarchical layers, which is why it is referred to as “deep learning”.

Deep learning is based on the concept of artificial neural networks (ANN), that is, systems inspired by the human brain’s structure (Goodfellow et al., 2016, p. 165). They are termed networks because they typically comprise a large number of interconnected nodes, referred to as neurons, distributed across various layers (Goodfellow et al., 2016, p. 164). **Figure 2** illustrates the simplified structure of an ANN.

The input layer receives the initial information (for example, raw data), which is then transmitted to the subsequent layer. Each node or neuron processes this input through a simple computation and forwards its output to the next layer of neurons. The artificial network “learns” by modifying the parameters of the connections between individual neurons, known as weights, based on experience and performance. This adjustment is made using a function that calculates the discrepancy between the actual output and the desired output, then alters the weights to minimize this error (LeCun et al., 2015, p. 436) This process, known as back-propagation, was first detailed in an article by Rumelhart et al. (1986, p. 533) and enables the program to self-organize and refine its internal structure. The layers between the input and output are called “hidden layers” because their activities are not directly observable from the input data. Instead, the model must determine which patterns are significant for explaining

the input data (Goodfellow et al., 2016, p. 6).

There are different types of artificial neural networks. Two of these are recurrent neural networks (RNN) and convolutional neural networks (CNN). While RNNs specialize in the processing of sequences, such as text (Graves, 2012, p. 1), the strength of CNNs lies in the processing of raster information, such as images (Goodfellow et al., 2016, p. 367). A weak point of the RNN and CNN architectures is the limited context window, i.e., the amount of information that the system can store over a longer period of time (Hochreiter et al., 2001, p. 11). To mitigate this vulnerability, Hochreiter and Schmidhuber (1997) developed a novel RNN architecture called Long Short-Term Memory (LSTM). Their customization made it possible to solve larger and more complex tasks that were not feasible with the standard RNN architecture (Hochreiter & Schmidhuber, 1997, p. 2).

However, the processing of long sequences remained a major challenge. This became particularly apparent in the area of machine translation, as it was difficult for RNNs to recognize word dependencies over long distances in sentences or paragraphs (Bahdanau et al., 2014, p. 6; Kim et al., 2017, p. 2). One idea was not to encode the entire sequence, but rather to focus on individual sections. The so-called “attention mechanism” allows the model to dynamically focus on different parts of the input (Kim et al., 2017, p. 2). To simplify, this can be compared to reading a scientific article and stumbling over a difficult section. Instead of trying to comprehend the entire content, it is helpful to focus on individual words or to look back at the previous sections. In this way, attention is shifted to the key aspects, helping to break down the section or sequence into smaller, more manageable parts and to understand how each part connects to the others.




Figure 2: Simplified structure of an Artificial Neural Network (ANN)

3.1.2. Emergence of Attention Mechanisms and Transformers

The attention mechanism described above was very effective, especially for machine translation tasks, and is therefore regularly used in deep learning models (Kim et al., 2017, p. 1). However, the sequential nature of RNNs remained the limiting factor in terms of parallelization and scalability of these models (Peng et al., 2023, p. 1). In 2017, Vaswani et al. proposed a new architecture in their article “Attention Is All You Need”. Called the “Transformer”, this new model is based solely on the attention mechanism and no longer on an RNN or CNN structure. Its strength lies in its ability to reference infinitely long sequences without the previous limitations of RNNs (Vaswani et al., 2017, p. 5), as visualized in **Figure 3**.

If the model is given the task to answer a question, it can refer to the entire text with the help of the attention mechanism. This is particularly helpful when textual relationships only become apparent upon viewing an entire paragraph. RNNs, with their limited context window, were restricted in this regard. In the example, an RNN model would “forget” that the question referred to a cat.

In 2015, the research group OpenAI was formed with the goal of advancing in the field of AI (OpenAI, 2015). One of the first articles published by OpenAI’s researchers related to the Transformer architecture was “Improving Language Understanding by Generative Pre-Training” (2018). In the article, Alec Radford and Karthik Narasimhan focus on enhancing natural language processing (NLP) via a semi-supervised learning approach that combines unsupervised pre-training and supervised fine-tuning (2018, p. 1). The authors propose a novel approach that utilizes large unlabeled text corpora for pre-training a language model, followed by task-specific fine-tuning. Their approach aims to overcome the limitations of supervised models, which often require extensive labeled data that is scarce or expensive to obtain.

In their model, Radford and Narasimhan (2018, p. 2) em-

ploy the Transformer architecture due to its efficient handling of long-term dependencies in text sequences. Their training method consists of two stages: In the pre-training phase, they use a language modeling objective on an unsupervised corpus, establishing the initial parameters of the model. Language modeling is the process of training a model to predict the next word based on previous words in a sentence (Voita et al., 2019, p. 3). **Figure 4** shows a simplified example using a Google search, where the model suggests the next word based on its predicted probability.

Applying this principle to training data, a set of 7,000 unpublished books, the model was able to recognize correlations and gain knowledge not only about individual sentences but also about the nature of language in general (Radford & Narasimhan, 2018, p. 8). In the fine-tuning stage, the learned correlations and knowledge were adapted to specific language understanding tasks such as question answering, common sense reasoning, semantic similarity analysis, and text classification. Subsequently, they used different tests to measure the performance in these categories, achieving state-of-the-art results in 9 of 12 categories (Radford & Narasimhan, 2018, p. 8).

3.1.3. Advancements in GPTs and Foundation Models

Radford and Narasimhan’s (2018) work laid the basis for further development of the GPT architecture. In an article published in 2019, a next-generation GPT model, GPT-2, was introduced (Radford et al., 2019). Instead of using curated datasets, they based their training on publicly available data from the internet. With the intention of creating a large and diverse corpus of natural language text covering as many domains as possible, they gathered over 8 million documents, totaling 40 GB of data (Radford et al., 2019, p. 3). When testing for natural language processing tasks, such as question answering, translations, reading comprehension, or summarization, they observed the model’s capability to learn these tasks without explicit supervision (Radford et al.,

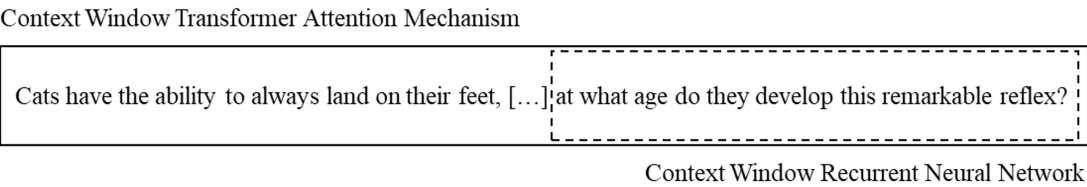


Figure 3: Comparison of context windows




Figure 4: Google Search: Example of word predictions

2019, p. 1). This can be seen as an significant breakthrough, as previous models required substantial amounts of labeled and curated data to learn effectively (Radford & Narasimhan, 2018, p. 1). During their tests, they experimented with different model sizes and showed that performance improved with increasing model capacity. Their largest model, with 1.5 billion parameters, achieved state-of-the-art results in multiple settings (Radford et al., 2019, p. 10).

The potential for performance gains through increased model capacity was further explored by Brown et al. (2020). Their article introduced the GPT-3 model, containing 175 billion parameters and representing a significant advancement over previous models by a factor of 10. Confirming the researchers' hypotheses, the model outperformed its predecessor, GPT-2, and even rivaled the performance of state-of-the-art fine-tuned systems (Brown et al., 2020, p. 9). They found that additional computational power directly correlated with increased performance, laying the groundwork for further research. Moreover, since their work focused on creating a task-agnostic model, they anticipated that future fine-tuning would further enhance their model's performance (Brown et al., 2020, p. 2).

Further developments occurred over the next years, advancing the model and introducing fine-tuning to boost performance (Ouyang et al., 2022, p. 1). This culminated in OpenAI's release of the now well-known ChatGPT (OpenAI, 2022). With its unexpected capability to generate text indistinguishable from human writing and engage in believable human-like conversations, it caught many researchers off guard (Dwivedi et al., 2023, p. 4). Technologically, the advances stem from two primary aspects. Firstly, an updated GPT-3 model with a modified dataset serves as the foundation for ChatGPT (GPT-3.5). Secondly, a technique called Reinforcement Learning from Human Feedback (RLHF) was used to fine-tune the model. RLHF can be defined as a method where the model is improved based on human eval-

uations of its outputs, guiding it toward desired behaviors and responses (Ziegler et al., 2019, p. 1). The process utilizes a reward model that is trained using feedback from human labelers, who evaluate and rank the model's answers from best to worst. This reward model is then used within the main model, employing reinforcement learning for self-optimization. This significantly enhances the model's ability to provide answers aligned with human preferences and consider previous input in conversational settings (Feuerriegel et al., 2023, p. 4). Alongside the technological leap, OpenAI's decision to grant unrestricted public access to ChatGPT and its simplified user interface fueled its explosive worldwide adoption (Feuerriegel et al., 2023, p. 5). **Figure 5** shows the landing page of ChatGPT 3.5, offering examples of questions and a prompt bar.

In March 2023, just five months after ChatGPT's release, OpenAI unveiled the next generation of their GPT model (OpenAI, 2023). Unlike previous models focused solely on textual input, GPT-4 could process both text and images as input, generating textual outputs. While detailed specifications remain limited, its performance exceeds that of its predecessor, GPT-3.5, in numerous areas, achieving human-level results on multiple benchmarks (OpenAI, 2023, p. 6). As prior research demonstrated a correlation between performance and computational power, anonymous reports suggest the model may possess 1.76 trillion parameters (Bastian, 2023). While this information should be treated with caution, it gives an indication of the rapid technological development in the area of Generative AI. From 1.5 billion parameters in 2019, over 175 billion in 2020 to potentially 1.76 trillion in 2023.

With the advances in the field of AI, a shift in the development of AI models has emerged. While previous models were trained on labeled data and fine-tuned to specific tasks, newer models such as GPT-3 and 4 differ significantly in two areas. Firstly, instead of labeled data, they are trained

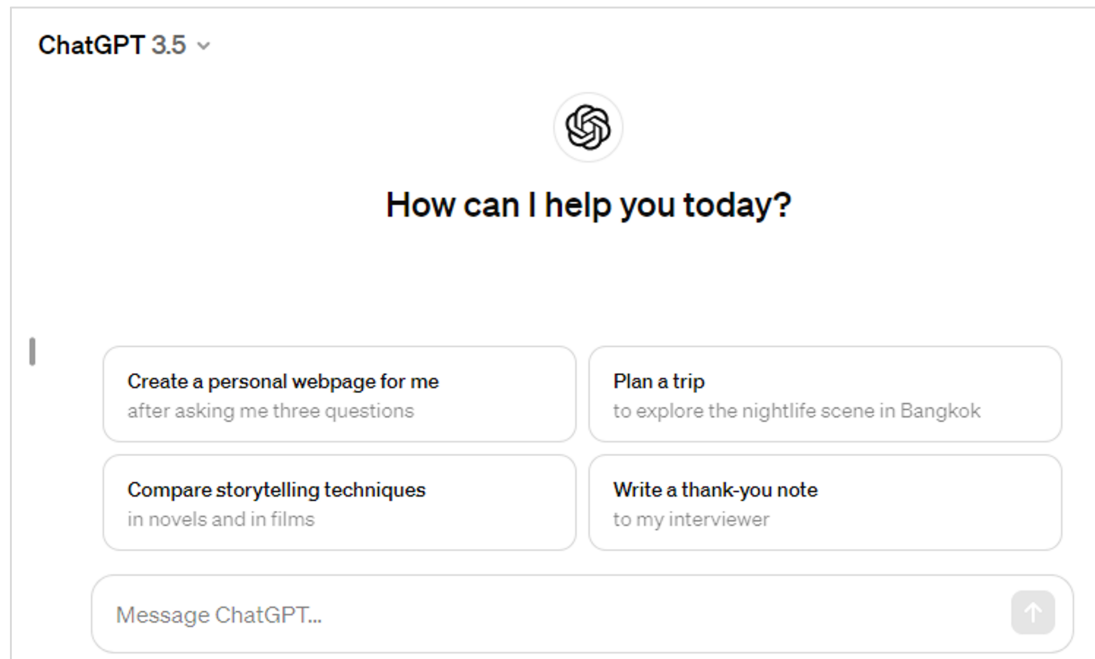


Figure 5: ChatGPT Landing Page.

on broad, unlabeled information using self-supervised learning. Secondly, the sheer scale allows the models to perform tasks they have never been explicitly trained on (Brown et al., 2020, p. 9). This general understanding of language and other patterns, fostered by the increase in available computational power, led to the need to create a new class for these models. In their scientific article, Bommasani et al. (2021) therefore proposed the term “foundation model” to encapsulate these characteristics and describe the emergence and application of these models in detail. According to the authors, foundation models build on the idea of large language models, which gain an understanding of natural language based on textual datasets (Brown et al., 2020, p. 9). Expanding on this concept, foundation models are multimodal – besides text, they are able to process data such as images, videos, 3D signals, and others. As with large language models, their vast datasets allow them to gain a general understanding of the underlying structures and properties applicable to different downstream tasks (Bommasani et al., 2021, p. 6). A prominent example is GPT-4, with its capability to provide answers based not only on text but also on images (OpenAI, 2023).

3.1.4. Generative AI: Scope and Capabilities

In the previous section, various development steps and technological innovations have been outlined, from simple ML algorithms to deep learning techniques, culminating in GPTs, large language and foundation models. One feature that was mentioned in connection with GPTs is their ability to generate new content. While the concept of GPTs has its origin in the publicly available work by Vaswani et al. (2017) on the Transformer architecture, the term itself was primarily coined by OpenAI’s advances in the field (in particular

with GPT-3 and 4). Google, Meta, and other companies have since published their own models based on the Transformer model (for example Google’s Gemini or Meta’s Llama 2). To maintain a neutral perspective, the term “Generative AI” is therefore used throughout this thesis.

Generative AI, from a technical standpoint, is rooted in generative modeling, which seeks to understand the joint probability distribution $P(X, Y)$ where X represents the data and Y denotes the labels. The objective is to grasp how data is generated in order to create new data points (Ng & Jordan, 2001, p. 1). This approach differs from discriminative modeling, which focuses on modeling the conditional probability $P(Y, X)$, representing the probability of the label Y given the input data X (Ng & Jordan, 2001, p. 1). This differentiation is crucial from an application standpoint: Generative AI models enable the generation of new data instances based on observed probability distributions within a given dataset (Goodfellow et al., 2016, p. 716), whereas discriminative modeling exhibits superior performance in classification tasks (Ng & Jordan, 2001, p. 1).

3.2. AI Value Creation and Business Impact

Over the last few decades, a large number of scientific articles have focused on the influence of AI on companies. The question of how and in which areas AI can be used and what prerequisites companies need to fulfil in order to do this successfully is a recurring one. Selected theoretical concepts will therefore be presented and described in the following section. Caner and Bhatti’s (2020, p. 182) conceptual framework can be used as a starting point. Based on an extensive literature review, the authors propose six perspectives from which the strategic dimension of AI can be viewed:

1. Capabilities and Limitations of AI
2. Business Functions and AI
3. Tasks, Jobs, and Intelligence
4. Economy and AI
5. AI and Law, Regulations, Governance
6. Industries and AI

This thesis focuses on the first three aspects, as they are especially useful for understanding how AI impacts individual companies. The first aspect covers the **capabilities and limitations of AI**, as discussed in Chapter 3.1. It explores the technical aspects of AI, including its various applications, technologies, and limitations. Understanding AI's strengths and weaknesses is crucial in a business context to set realistic expectations and identify practical deployment options (Davenport & Ronanki, 2018, p. 110). While some constraints identified by Caner and Bhatti (2020), like data labeling and learning generalizability, have been addressed with technological progress (refer to OpenAI, 2023), challenges such as biases and lack of explainability persist.

The aim of the second dimension **business functions and AI** is to categorize the fields of application of AI. Caner and Bhatti base this on an IBM survey and a study by Davenport and Ronanki (2018) (Caner & Bhatti, 2020, p. 186). They differentiate between two subject areas: Firstly, the areas in which AI is used, e.g., in internal processes or customer services. Secondly, the way in which AI is used, e.g., to automate entire processes or gain new insights from data.

The third dimension, **tasks, jobs, and intelligence**, describes the different stages in which AI can be used. Based on Rao's (2017) framework, a distinction is made between three levels. Firstly, "Assisted Intelligence" includes all processes that are conducted in a conventional manner and are supported by AI. Secondly, "Augmented Intelligence" refers to areas where AI takes over a large part of the value creation (the areas mentioned are, for example, automatic translation and automatic analysis of legal documents). Thirdly, "Autonomous Intelligence" comprises those processes that can be conducted completely without human intervention in the future.

3.2.1. AI in Strategic Business Context

The progress that has been made in the field of artificial intelligence repeatedly raises the question of how this technology can be strategically implemented and utilized. In particular, the question of the sources of value creation is being researched in depth (Borges et al., 2021; Kitsios & Kamariotou, 2021; Trunk et al., 2020). This aligns with the thesis's objective, namely, to investigate the extent to which the existing concepts can be applied to Generative AI. Borges et al. (2021) extensive literature review provides the first insights into this area. Based on 41 studies, they define four sources of value creation from AI:

1. Decision Support
2. Customer and Employee Engagement
3. Automation
4. New products and services

Decision support refers to the ability of AI to support humans in strategic and operational business decisions (Borges et al., 2021, p. 11). This occurs, for example, when deep learning techniques are used to detect patterns in data, subsequently guiding the decision-making processes. Consequently, decisions can be executed more swiftly and with greater reliability (Borges et al., 2021, p. 12). However, scientists highlight the necessity for further research in this rapidly evolving field, particularly concerning the interaction between humans and AI and its impact on organizational performance (Lichtenthaller, 2019, p. 8). This thesis presents an opportunity to examine these dynamics, especially in the context of Generative AI. It will be insightful to determine whether and how the technology can optimize decision-making processes and thereby contribute to value creation.

The second identified source of value creation is called **customer and employee engagement**. The aim is to use AI to improve the customer experience and internally to attract employees to the new technology. Although the article analyses various academic papers, the authors note that further research is needed as some examples cannot be generalized. Generative AI, with the ability to create content that is indistinguishable from human input (Feuerriegel et al., 2023, p. 1), could offer new opportunities in this context.

Automation represents the next source of value creation, allowing costs to be reduced and efficiency to be increased through the strategic use of AI technologies. The authors argue that a competitive advantage can also be created if processes can be automated more quickly than is possible for competitors (Borges et al., 2021, p. 12).

The fourth source of value creation, **new products and services**, deals with the ability to generate new business ideas through AI. Borges et al. (2021) suggest AI's potential in driving innovation and creating new products and services, however they could only identify and provide limited empirical evidence. Consequently, the authors argue that further research is needed to understand how AI can be used strategically to create new products and solutions (Borges et al., 2021, p. 12).

The sources of value creation listed by Caner and Bhatti as well as Borges et al. can also be found in similar form in other literature reviews and scientific articles. For example, in their literature review of 81 articles in the context of AI and business strategy, Kitsios and Kamariotou (2021) identified "AI and Machine Learning in organizations", "AI, knowledge management and decision-making" and "AI, service innovation and value" as sources of value creation.

An area that was not mentioned in earlier sources and which they discovered was the **alignment of AI tools and Information Technology (IT) with organizational strategy**.

According to the authors, the company's IT strategy should be closely aligned with the business strategy. They suggest developing a separate digital (business) strategy to take account of the increasing importance of IT (Kitsios & Kamariotou, 2021, p. 6). With regard to AI, the authors note that it can now perform cognitive tasks, i.e., processing new information, recalling it from memory and using it in communication with humans. One obstacle to the implementation of these technologies however is not the technology itself, but human perceptions, company processes and cultural barriers (Kitsios & Kamariotou, 2021, p. 7).

3.2.2. AI and Business Capabilities

When introducing AI, the question regularly arises as to what capabilities companies need. These requirements have been analyzed in detail in various scientific articles, each with a different focus area. Based on a comprehensive literature review, Reim et al. (2020) approach the topic from a business model perspective. They derive a four step model to provide guidance to companies (Reim et al., 2020, p. 180): In their view, companies need to develop a good understanding of AI, have a thorough understanding of their business model and possible areas for innovation, must invest in the needed capabilities to support AI implementation and aim for a high level of organizational acceptance. From their perspective, the initial steps will likely be driven by the upper management, laying the conceptual foundation for the introduction of AI and analyzing the required capabilities (Reim et al., 2020, p. 187).

Based on previous research on business process thinking and knowledge management, Uren and Edwards (2023) propose the following categories: People, Processes, Technology and Data (Uren & Edwards, 2023, p. 7). In their view, data in particular plays a crucial role as it provides the foundation for the successful adoption of AI (Uren & Edwards, 2023, p. 6). They add that a good understanding of the interplay of these capabilities is key to evaluate the impact of AI. Slightly different capabilities were identified by Brock and Wangenheim (2019). Instead of People and Processes, they propose Strategy and Security as key capabilities alongside Technology and Data (Brock & Wangenheim, 2019, p. 116).

Expanding upon these studies, further research was undertaken to explore and substantiate the needed capabilities. Focusing on organizational capabilities, Weber et al. (2023) identified four dimensions: AI project planning, co-development of AI systems, data management and AI model lifecycle management. Regarding AI project planning, they emphasized the need to identify and evaluate suitable use cases, as 76% of companies struggle to scale AI successfully across their organizations (Reilly et al., 2019, p. 3). In this context, understanding the specific characteristics and limitations of AI models was mentioned as an important prerequisite to better understand the possibilities of AI and to manage expectations (Weber et al., 2023, p. 1555). As pointed out by the researchers: "There is much value in professionalizing and democratizing the process of use case generation to collect fruitful AI use cases" (Weber et al., 2023, p. 1556). They

state that each use case should have a clear value proposition to enable the company to prioritize and pursue the ones with the highest value creation potential.

Another important aspect that was frequently mentioned is the collaboration between technical and business functions (Uren & Edwards, 2023, p. 7). It is pointed out that data scientists need to work closely together with business experts to drive AI projects successfully (Akkiraju et al., 2018, p. 6). Including domain experts into the process plays an essential role as AI solutions can have a major impact on existing work structures. Co-development between IT professionals, data scientists and domain experts can help to reduce potential fear of AI systems and help to drive the implementation more effectively (Weber et al., 2023, p. 1556).

Mentioned in various scientific papers, data management and capabilities seem to be a key factor for successful AI adoption (Brock & Wangenheim, 2019; Uren & Edwards, 2023; Weber et al., 2023). The authors mention that many use cases do not make it into productive implementation due to insufficient data quality or unsuitable data structure (Weber et al., 2023, p. 1557). Implementing a clear governance process for the collection, storage and curation of data is therefore a vital step to ensure the successful application of these use cases in a real-world setting (Uren & Edwards, 2023, p. 6). In this context, it is important for organizations to not only assess their existing data management strategies but also to anticipate and plan for the evolving demands of future AI solutions (Weber et al., 2023, p. 1557).

Implementing new AI solutions raises questions about their long-term management in large corporations. Weber et al. (2023) highlight that this issue is particularly relevant given the iterative development process of AI models, involving multiple tests and adjustment cycles. Coordinating these releases is critical, especially when AI models are implemented across various organizational units, each potentially requiring minor or major adjustments. Consequently, organizations must ensure that AI solutions are integrated into their IT landscape and aligned with their business processes. The ability to manage these steps in the AI models' lifecycle is thus considered a crucial factor in successfully introducing AI into an organization (Weber et al., 2023, p. 1557).

Across many publications, special attention was given to the people dimension. Analyzing capabilities for digital transformation, Blanka et al. (2022) identified employee transformation competency to be crucial for driving digitalization in companies. They observed that employees as well as managers need to have a basic understanding of the potential of digitalization in order to effectively leverage it for organizational innovation (Blanka et al., 2022, p. 10). To develop these skills further and to be ready for the introduction of AI, broad training measures are required (Lee et al., 2019, p. 8). According to the authors, AI experts are rare. Offering training to the workforce could therefore help to improve understanding and support the digital transformation process. This also seems to be applicable for Generative AI (Feuerriegel et al., 2023, p. 6). Expanding the knowledge about the technology in the organization could

help identifying use cases and increase the overall readiness to integrate Generative AI into everyday business processes. To overcome organizational resistance and successfully scale AI applications, Lee et al. (2019) propose the following five step process model, as displayed in **Figure 6**.

Starting with executive pilot projects, the authors argue that companies should focus on smaller applications first. This supports the goal to acquire knowledge and to celebrate success stories early on. Based on these success stories, use cases can then be scaled to larger projects. This is accompanied by the formation of dedicated AI-teams which drive the topics forward. The authors acknowledge that this is likely not possible for smaller companies or startups which should consider support from 3rd party providers. Besides the broad AI-training, the development of an AI-strategy is highlighted as a crucial aspect. While Lee et al. (2019, p. 8) focus on the importance of data in this context, it can be argued that a more holistic view is required as outlined by Brock and Wangenheim (2019) and Uren and Edwards (2023). Finally, Lee et al. (2019) highlight the importance of communication to all stakeholders. Due the novel character of AI and corresponding uncertainties, they point out that communication about the advantages and potential challenges is needed. They add that this could also help to find new, innovative ideas for AI when discussing use cases with customers (2019, p. 9).

4. Methodology

The research design of this thesis follows an exploratory case study approach as outlined by Yin (2018) with a clear methodological path, which is based on a comprehensive literature review as described in Chapter 3. For background information on the technology and related scientific research, two major streams have been followed. Firstly, the development of artificial intelligence, starting with the first conceptual ideas in the 1950s up to the latest advances in the field of Generative AI. Secondly, the review of scientific research on artificial intelligence and its value creation opportunities. Given the novelty of Generative AI, most available concepts and models have been written and published before its introduction, dealing with the available AI at its time. The following case study therefore offers a great opportunity to add new insights to the research field.

Supported by the literature review, an exploratory case study design is applied, an approach especially feasible for phenomenon which have not yet been comprehensively researched (Yin, 2018, p. 15). According to Yin, a case study is an empirical method to investigate a phenomenon or “case” from multiple angles to create an in-depth understanding. It is especially useful when the phenomenon can only be analyzed within its real-world context and its boundaries are not clearly visible. The goal is therefore to create an in-depth understanding of the subject and its current application which should allow for analytic generalization (Yin, 2018, p. 21). Doing so, the case study follows a single-case study design

which allows for a focused study of the selected subject and rich data collection (Yin, 2018, p. 49).

4.1. Case Study Design

In the search for a suitable object of investigation, two essential questions were answered. First, does the case provide a good basis for exploring the value creation potential of Generative AI? And secondly, is it possible for the author to gain valuable insights for the research within the scope of his possibilities? By defining the company Siemens as the subject of the study, both questions could be answered in the affirmative. On the one hand, the company with its multinational setup and different business areas offers ideal conditions to gather diverse perspectives on Generative AI. On the other hand, as an employee of the company, the author is able to gain insights that are not available to external parties. The latter circumstance in particular promises to provide valuable insights into the company and its approach to new technologies. Moreover, it offers a unique opportunity to validate scientific theories and concepts using a real-life example.

Founded in 1847, Siemens is a multinational corporation headquartered in Munich. Siemens sees itself as a technology company with the mission to “provide technologies that improve quality of life and create lasting value for society” (Siemens, 2024). The company focuses on four core businesses: Digital Industries, Smart Infrastructure, Mobility and Siemens Healthineers. Further businesses include Siemens Financial Services and its portfolio companies Innomotics and Siemens Logistics (Siemens, 2023, p. 85). Corporate functions include among others Strategy, Compliance, Supply Chain Management and Global Business Services as well as IT, Human Resources, Controlling and Communications. In fiscal year 2023, ending on September 30, 2023, Siemens generated revenue of € 77.8 billion from continuing operations with a corresponding net income of € 8.5 billion. As of September 30, 2023, around 320,000 people were employed by the company worldwide (Siemens, 2024).

The case study was able to collect information from almost all areas of the company. Only information from Siemens Healthineers had to be omitted, as it is managed as an independent company within the Siemens Group. In addition to interviewing managers and employees from Digital Industries, Smart Infrastructure and Mobility, participants were recruited from various business functions, including Human Resources, Global Business Services, IT, Strategy, Supply Chain Management, R&D as well as Technology.

4.2. Data Collection

Semi-structured interviews with key informants form the core of the study. These interviews provide helpful qualitative data and insights from various perspectives. The interview partners were selected specifically in order to obtain a broad view of the research subject. This “triangulation of subjects”, as outlined by Rubin and Rubin (2012, p. 63), aims to bring together different ideas, perspectives, and opinions. In addition, when selecting the contact persons, attention




Figure 6: Five Step Process Model, based on Lee et al. (2019)

was paid to their background knowledge of the technology in order to generate high-quality input for the study (Rubin & Rubin, 2012). The interview partners were from various areas, roles and responsibilities, including CFOs, CIOs, Program Managers for Generative AI, Product Owners, Domain and Technology Experts as well as IT Professionals and Innovation Managers (see Table 1).

Interviews lasted between 24 minutes and 1 hour and 15 minutes. Due to separate locations and time zones of the audience, all interviews were held virtually via Microsoft Teams, using the latest technology for recording and transcription. To allow for an open exchange and unfiltered perspectives, the confidentiality of the meetings was agreed upon and all insights will only be used in an anonymous way.

A total of 23 interviews were conducted in the period from December 2023 to March 2024 with an average length of 43 minutes. MAXQDA, an external software designed for computer-assisted qualitative data analysis, was used to analyze the interviews. During the process, the transcripts were anonymized, cleaned, and organized. It was found that many transcripts contained erroneous sections and thus required manual correction of words. This was due to the fact that the software used for transcription, Microsoft Teams, did not correctly recognize the interviews that were conducted in German. Furthermore, due to the specific topic and many abbreviations used within Siemens, many technical terms were not correctly recognized by the software. Manual coding was therefore conducted with the help of the transcripts and video recordings. New interviews were arranged until it was ensured that all aspects were sufficiently covered, and no blind spots remained. Reaching this 'saturation point' which was first described by Glaser and Strauss (1967, p. 61) contributes to the depth and completeness of the research findings.

A guideline was drawn up for the interviews in order to give the discussions a clear structure and to facilitate the subsequent evaluation. Table 2 shows the general structure and topics of the interview guide. In addition to the guiding questions, the interview partners were given space to address additional aspects and explain them in detail. This ensured that findings and topics not yet covered by questions were also recorded (Myers, 2020, p. 149). The interview guide was successively revised after each interview to take new aspects of the data collection into account. This helped to verify aspects raised by interview partners in subsequent interviews.

Besides the semi-structured interviews, other sources were used to supplement and enrich the case study. From

the six sources of evidence, as outlined by Yin (2018, p. 113), the following are included in the thesis

- Interviews

Interviews, as previously discussed, are conducted to gather valuable insights, perspectives, and explanations from individuals who possess key information. This process follows a structured interview guide containing predefined topics and questions. However, it also allows for flexibility to explore alternative inquiries and different areas of focus (Rubin & Rubin, 2012, p. 31).

- Documentation

The data collection process for the case study involves systematically searching for relevant literature. This serves as a crucial component for future reviews and data triangulation. It comprises several types of data, including press releases, internal presentations, meeting minutes, reports, news articles, wikis, and information stored on SharePoint.

- Direct Observations

Direct observations provide valuable additions to the mentioned interviews. These observations encompass various activities within the context of the research phenomenon, such as conferences, meetings, recorded sessions, training courses, and other relevant events. Furthermore, internal communities dedicated to the research topic offer insights into the phenomenon and serve as an additional source of information.

- Participant Observation

The author of this thesis does not merely act as a passive observer but actively engages in exploring Generative AI within the case study's context. In his current role within the company, he delivers presentations on the topic, participates in conferences and meetings related to Generative AI, and engages in discussions with colleagues regarding the technology's ongoing development. As described by Yin (2018, p. 124), this active involvement presents both unique opportunities and risks, which have been carefully evaluated for their potential impact on the case study. Firstly, as an active participant, the author can gain practical knowledge in addition to theoretical insights from his firsthand experiences. Secondly, it enables the author to partici-

Table 1: Overview of Interviews

No	Role	Organization Type	Length	Date
1	Program Manager	Human Resources	55 mins	Dec 23
2	Program Manager	Mobility	52 mins	Dec 23
3	Product Owner	IT	44 mins	Dec 23
4	Domain Expert	Supply Chain Management	55 mins	Dec 23
5	IT Professional	IT	51 mins	Jan 24
6	Product Owner	R&D	54 mins	Jan 24
7	Domain Expert	Supply Chain Management	54 mins	Jan 24
8	CFO	Mobility	45 mins	Jan 24
9	Technology Expert	IT	27 mins	Jan 24
10	Program Manager	Global Business Services	42 mins	Jan 24
11	Technology Expert	IT	29 mins	Jan 24
12	Domain Expert	Mobility	1h 15 mins	Jan 24
13	Domain Expert	Mobility	52 mins	Jan 24
14	Program Manager	Technology	53 mins	Jan 24
15	CIO	IT	24 mins	Jan 24
16	Head of Technology	Digital Industries	33 mins	Jan 24
17	Domain Expert	Digital Industries	30 mins	Feb 24
18	Technology Expert	Technology	35 mins	Feb 24
19	Domain Expert	Smart Infrastructure	26 mins	Feb 24
20	Strategy Manager	Strategy	30 mins	Feb 24
21	Program Manager	Smart Infrastructure	40 mins	Feb 24
22	Program Manager	Smart Infrastructure	54 mins	Feb 24
23	Innovation Manager	Digital Industries	32 mins	Mar 24

pate in events and establish connections with individuals who might not have been included in the study otherwise. However, active involvement also poses significant challenges that require close monitoring (Yin, 2018, p. 124). Particularly, unintentional bias in supporting the research subject could lead to a subjective and skewed perspective. The author acknowledges this trade-off and takes utmost care to ensure an objective analysis.

The underlying goal of any research is to thrive for the highest possible quality. Its design should follow a logical sequence which allows other researches to analyze and verify its findings (Yin, 2018, p. 42). During the case study, the design quality was therefore measured and verified using three common tests. These are construct validity, external validity and reliability (Yin, 2018, p. 42). According to Yin (2018), construct validity aims at finding the correct measures to assess the underlying theoretical construct (Yin, 2018, p. 43). Using multiple sources, establishing a chain of evidence and reviews of the case study helped to increase the construct validity. The second test examines the degree to which the re-

sults of the study can be generalized to a different or broader context (Yin, 2018, p. 45). Grounding the research questions in appropriate theory, supports the external validity of the study. Ensuring reliability of the findings and conclusions, represent the final test (Yin, 2018, p. 46). To fulfil these requirements, the case study follows a clear approach and provides comprehensive documentation, allowing other researchers to potentially replicate the results.

4.3. Analytical Approach

As outlined above, the analysis of the data follows a six step approach of Braun and Clarke (2006, p. 15). This allows for a deep understanding of the gathered information and underlying patterns and provides a clear structure during the analysis. The first step involves the familiarization with the data. Structuring the entire data set, reading through the interviews and extracting first ideas and patterns, helps to lay the foundation for the in-depth analysis (2006, p. 16). In a second step, initial codes are generated. These codes summarize the meaning of a section, phrase or word and represent the building blocks for the analysis. They help to summarize

Table 2: Interview Guide

Introduction	
	Brief introduction and purpose of the study
	Clarification of confidentiality and anonymity
Part 1: Background Information	
	Walkthrough of professional background and experience with AI
	Current involvement in Generative AI
Part 2: Strategic Understanding of Generative AI	
	Comparison of Generative AI with previous AI models
	Functionalities and limitations of Generative AI and its relevance for Siemens
Part 3: Value Creation Opportunities	
	Use Cases - insights, development and value add for the company
	Application areas with the highest potential value creation
	Combination of Generative AI with other technologies
	Outlook on future value creation opportunities
Part 4: Expanding AI within Siemens	
	Requirements for successful implementation and scaling across Siemens
	Potential challenges and barriers (technical, organizational, etc.)
Part 5: Ethical and Security Considerations	
	Ethical considerations when introducing Generative AI
	Cyber and Data Security
Conclusion & Wrap-Up	

and classify the data and provide more transparency over the findings (2006, p. 18). The third step deals with the identification of themes. Based on the identified codes, patterns are extracted and merged to larger groups with the same meaning. These themes help to better understand ideas across the data set. Starting with the analysis, an inductive approach was followed allowing themes to emerge from the data itself (Myers, 2020, p. 210). In the fourth step, the identified themes are reviewed to ensure consistency and quality (Braun & Clarke, 2006, p. 20). This step included going back to the initial codes and verifying their validity and fit to the selected category as well as revealing potential biases and misunderstandings. In addition, the analysis as a whole was analyzed for ambiguities and overlaps. This proved to be a valuable step, as familiar patterns from observations in the company and aspects from the literature emerged. Following a deductive approach, these were considered as follows: In places where the data showed strong similarities with existing concepts, the terms were taken from the literature. For new topics, own terms were chosen. This approach made it possible to uncover previously unknown themes in relation to Generative AI and, at the same time, made it easier

to compare recurring themes with existing models and concepts (Myers, 2020, p. 210). Step five of Braun and Clarke's (2006) approach is closely linked to the previous one and deals with the definition and naming of themes. In this step, the emerging themes were once again refined to ensure that the underlying data was represented correctly (2006, p. 22). In this step, subcategories were identified for several themes which allowed for more transparency and a better understanding of the identified patterns. **Appendix 1** provides an overview of the themes and their occurrence in the interviews. The sixth and final step involves the presentation of the results. This will be done in Chapter 5, followed by a discussion in Chapter 6 which aim at providing a comprehensive insight into the findings and critical analysis of the data (2006, p. 23).

In total, 1,139 segments were identified in the raw data which were assigned to 429 codes. These were divided categories and subcategories, which were subsequently categorized into five main themes. **Appendix 2** shows the process from raw data, through open codes and categories to the main themes for a selected sample.

5. Results

The information for this case study was collected over 12 months from March 2023 to March 2024. During this period, the author participated in various meetings in which Generative AI aspects were presented and discussed. In addition, the author became a member of various groups actively involved in the introduction and dissemination of Generative AI in the company, regularly exchanged ideas with colleagues on new findings and participated in internal and external conferences. Furthermore, the author himself gave presentations and organized workshops to create awareness for Generative AI and to gain further insights. This was done regionally, in Asia Pacific, but also globally in management meetings with representatives from different business areas and functions. Moreover, in order to better understand the technological aspects of Generative AI, the author initiated a pilot project together with a team of data scientists (**Appendix 3**). The core of the case study form interviews with experts from various business areas as mentioned earlier. Finally, the research work is rounded off by the collection of internally and externally available documents.

5.1. Comparative Analysis: Generative AI vs. Previous Models

A key question that arises when it comes to Generative AI is what makes the technology so different and interesting for companies. Two key factors can be identified from the analysis: the technological innovation itself and its high usability compared to previous AI models. In terms of technology, three factors emerged from the data: improved textual understanding, learning speed and flexibility of the models. Interview partners pointed out that this makes it possible to realize use cases that were too complex with previous models

due to the effort required for adaptation and training. Illustrating this point, interview partner 11 (IP-11) stated:

“Foundation models allow us to solve many more use cases than before because we no longer have to collect data and train any model and that is [...] very disruptive for the AI world because that was always the number one problem before - where do I get my data?” (IP-11, Technology Expert)

The data also provides a clear picture with regard to usability. Interview partners state that interacting with artificial intelligence is now remarkably simple. The conversational user interface, which allows anyone to execute commands via prompts, is named as a key driver in this context. This empowers a wide range of employees to try out the tool without having to worry about complex control parameters or the underlying technology. As IP-7 puts it:

“The biggest added value of the interface itself is that you really only need text and speech to interact with [the] tool” (IP-7, Domain Expert)

The data suggests that this will open up new application possibilities in the future, as the models can now be used by employees without any technical background (IP-16, Head of Technology). In the future, for example, users could interact with a database for which SQL knowledge is still required today (IP-9, Technology Expert).

5.2. Value Creation Clusters

The collection and analysis of Generative AI use cases was an essential component of the case study. In addition to the pilots mentioned in the interviews, other internal sources of information were used to obtain a complete picture. This resulted in three perspectives from which potential value creation can be viewed. The first perspective involves categorizing the use cases according to the way in which the technology is used. The second perspective focuses on value creation drivers, i.e., use cases offering new value creation opportunities or improving existing structures. The third perspective divides the value creation opportunities in organizational units, allowing companies to identify key areas within their structures more easily. **Figure 7** illustrates the individual categories and subcategories.

Looking at value creation clusters, from the data collected, it became apparent that most of the use cases can be assigned to the **assistants**’ area. A further subdivision was therefore possible and allows for a more precise differentiation of use cases. It must be mentioned that during the case study, various approaches were identified at Siemens for the clustering of use cases. These were used as the basis for the initial categorization, revised and refined over the course of time. This resulted in a clear picture of the assignment of the individual use cases.

Within the first cluster, the area **data synthesis** deals with use cases that focus on extracting, analyzing, and gaining

knowledge from data. Generative AI is used to gain new insights that could not previously be achieved in this form or only with considerable additional effort. Examples include a Generative AI-supported tender analysis at Siemens Mobility to classify requirements and identify risks, the identification of common parts across multiple IT systems in Siemens Procurement or the analysis of employee comments as described in detail below.

Example: Employee Surveys

Siemens regularly asks employees about their satisfaction as part of a global campaign. This survey also includes free-text fields in which employees can write comments. With the help of Generative AI, a sentiment analysis can be carried out on all comments to determine whether the message was written in a positive, neutral, or negative tone. In this way, critical topics such as the risk of burnout can be identified without being explicitly mentioned in the text (IP-1, Program Manager). Consequently, Generative AI proves beneficial by allowing for an analysis without the need to set up and train a dedicated AI system specifically for sentiment detection.

Another potential use case that was regularly mentioned was the adaptation and training of foundation models to include Siemens specific content. While no advanced use cases were identified in the interviews, multiple research projects were mentioned, and many interviewees highlighted the potential for future applications. As expressed by one participant, the lower training effort compared to previous approaches promises a major advantage, as the models would already have a sufficient understanding of human language (IP-19, Domain Expert).

The second cluster, Generative AI-based **assistants**, contained by far the most application examples. The aim of these systems is to prepare and process existing information, making it readily available to the user. The subcategories are based on the complexity of the respective use case; factors include the type of data processed (static or dynamic), the number of systems connected, the query type (one-shot or staged), and the level of required human interaction. The necessary user knowledge also plays a role. While the simplest use cases require no prior knowledge, other systems are designed for IT professionals and data scientists.

With **information retrieval**, users are generally interested in finding information more quickly. These “factual chatbots” can access a variety of data sources, ranging from internal company wikis, guidelines, and circulars to department-specific knowledge databases. From a technological point of view, these systems use Retrieval Augmented Generation (RAG) in combination with a licensed large language model, approved for confidential internal information.

Use cases in the second subcategory **interplay with systems** require a close link between the LLM and the data source. The most common example cited was using Generative AI as a user interface to query internal databases. Here,

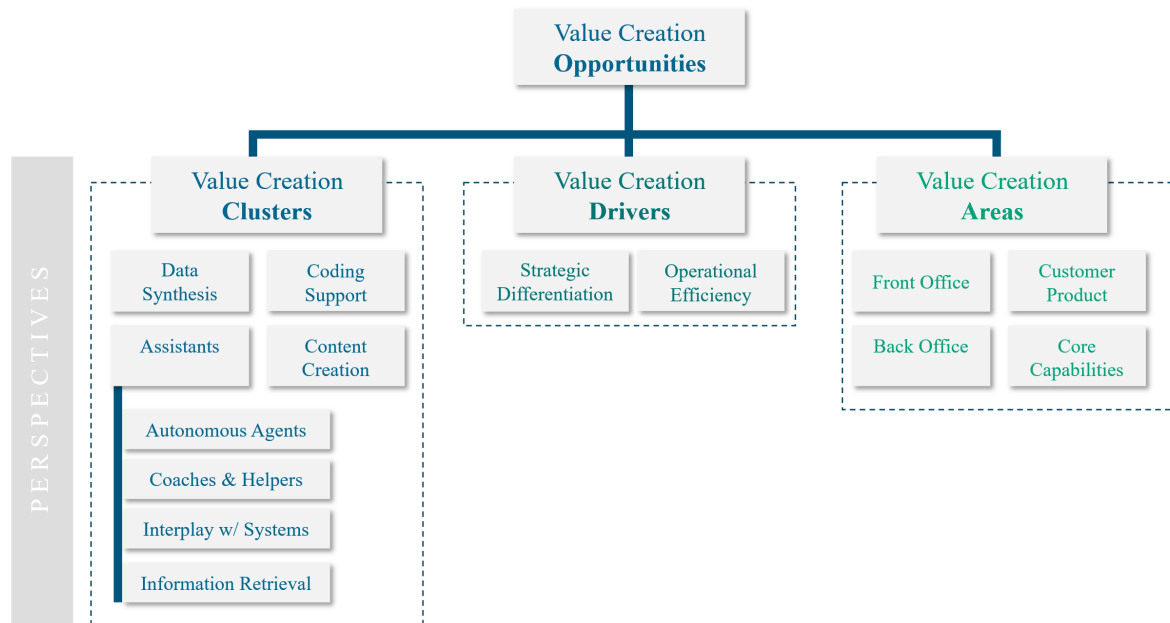


Figure 7: Value Creation Opportunities

Generative AI converts user commands in natural language into SQL commands to retrieve information. One interviewee emphasized that this could enable faster, more individualized queries in the future, decreasing the need to create new dashboards for end users (IP-4, Domain Expert). Other examples include combining Generative AI with Robotic Process Automation (IP-3, Product Owner), using Generative AI alongside SAP inputs (IP-7, Domain Expert), or evaluating application documents in the Siemens-wide job portal using Generative AI (IP-1, Program Manager).

The data revealed that the main focus of assistant systems is on use cases that actively support users in their tasks. These **coaches & helpers** accompany the user along the value chain in various process steps and influence the desired results. Added value is expected in two areas in particular. In the use of external tools such as Microsoft Copilot and in the development of internal, domain-specific assistants, such as a chatbot for service technicians in the area of Smart Infrastructure (IP-21, Program Manager). Further application examples include the development of chatbots to prepare and practice negotiations in purchasing (IP-7, Domain Expert), the use of Generative AI in interview processes (IP-9, Technology Expert) or support for process reviews as part of internal audits (IP-12, Domain Expert).

While the majority of the identified use cases envisage the above-mentioned assistance role for Generative AI, individual interview partners see the possibility of using the technology to **automate** entire processes. This usually includes the interaction of several software applications and systems, in the context of which Generative AI represents the interface to the user. Illustrating this point, an interviewee mentioned a Generative AI agent that makes it possible to classify customer inquiries automatically, extract the necessary informa-

tion, update the linked systems, and write a response to the customer (IP-10, Program Manager). It was pointed out that while this was already possible in theory, in reality the decisions would still be made by humans due to the maturity of the use case, the specific features of the technology and the associated risks.

Another large cluster comprises use cases that can be summarized as **coding support**. As a technology company, Siemens employs software in a wide variety of forms and areas. The ability of Generative AI to generate code in different programming languages was therefore seen as a great lever to increase productivity and improve existing software solutions. One example mentioned in the interviews and communicated externally by Siemens as a lighthouse project revolves around the Siemens Industrial Copilot, as described in detail below.

Example: Siemens Industrial Copilot

Siemens offers a variety of solutions for factory automation, including control systems for machines and processes. In this use case, Generative AI is deployed to write PLC (Programmable Logic Controller) code for these machines. A distinctive feature mentioned is the ability to translate and improve code from other programming languages. Additionally, the software leverages Generative AI to identify bugs and suggest solutions, significantly reducing programming effort and potential downtime in factories.

Another example of Generative AI used within Siemens' software solutions is its implementation in the Mendix low-code platform. In collaboration with external partner AWS, Siemens extended the program to further support users in the

development, validation, and optimization of Mendix applications. In addition to these two customer-facing solutions, many interviewees reported added value in developing internal applications. One example mentioned by a respondent was Generative AI assisting in the development of an automation script for a building's energy supply (IP-19, Domain Expert). A common theme from the interviews was that the coding environment's focus was less on individual use cases and more on transforming the way programming is done. Many respondents see great added value in the ability to generate and improve code, as well as in using assistants like GitHub Copilot (IP-6, Product Owner).

The last cluster can be described as **content creation**, where the focus is on generating new insights and data with the help of Generative AI. Common examples include its use in marketing for writing news articles or in human resources for generating job advertisements. In addition to these general applications, specific Siemens use cases were identified. For example, one interview partner reported the use of Generative AI as part of the innovation process to find new alternative materials for products (IP-7, Domain Expert). Another participant reported using Generative AI to translate 3D models of metal parts into code for milling machines (IP-18, Technology Expert).

In summary, the collected use cases indicate a broad range of applications for Generative AI. In addition to simple models for information retrieval, the technology also offers added value for complex internal processes as well as customer-facing applications. In this context, the identified value creation clusters provide a valuable starting point for better understanding the diverse potential of these use cases.

5.3. Value Creation Drivers

The data analysis revealed two main drivers of Generative AI use cases, as highlighted in **Figure 8**.




Figure 8: Value Creation Drivers

The first group can be broadly summarized as **strategic differentiation**. These use cases focus on gaining a unique competitive advantage by leveraging resources exclusive to the company, enabling differentiation from the competition. These resources include internal company data (such as customer information), process knowledge, and the IT infrastructure itself. This group also encompasses the development of new products or the expansion and application of Generative AI capabilities within existing products. Examples include the Industrial Copilot, which Siemens uses to strengthen its own software applications, and the integration of Generative AI into the Mendix low-code platform.

Internally, this can involve developing technology stacks and platforms to accelerate the rollout and distribution of

specific use cases. For instance, the Global Shared Services division developed a Generative AI platform that enables departments worldwide to adopt simple use cases without needing to invest in the technology themselves (IP-10, Program Manager). This creates synergy effects, as the internal solution requires development only once and can subsequently be rolled out across the entire organization.

The second driver is productivity improvement. The data and interviews revealed that the majority of current use cases fall into this category, less focused on unique solutions and more on leveraging the technology effectively. This includes making solutions broadly accessible to as many users as possible to improve **operational efficiency** on a large scale. Participants frequently mentioned faster information retrieval as one of the biggest levers. Others listed quality improvements, faster execution of processes, and risk minimization as productivity-boosting factors.

5.4. Value Creation Areas

One finding that emerged during the case study is the company's need to cluster potential use cases not only according to its technological background, but in particular according to company-specific segmentation. This was particularly noticeable when analyzing company presentations and participating in meetings on the possible areas of application of Generative AI. Depending on the focus, different perspectives can be favored in the various company divisions. **Table 3** shows three distinct views on use cases which are used within the company. The first view is built around functional areas within the company and allows interested parties to differentiate Generative AI use cases according to an established company-wide used segmentation. The second view is built around business capabilities which describe what businesses do to achieve a specific purpose. Driven by processes and not by organizational structures or technology, this view is often used in IT for strategic analysis. Lastly, clustering in opportunity areas provides a more generic view which is in particular suitable for gaining a faster understanding of Generative AI application areas. Originally proposed by Gartner (2023), Siemens adopted the model for its own organization. It distinguishes between front and back-office tasks, product, and services as well as core capabilities of a company.

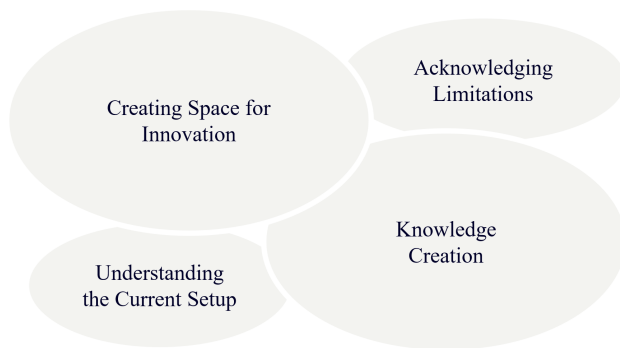
5.5. Exploration and Scaling

While focusing on the value creation opportunities of Generative AI, the case study identified several supplementary elements. Many interviews provided insights into the current state of Generative AI at Siemens and factors that should be considered for a successful company-wide rollout. Analyzing the data, two major patterns emerged:

Firstly, Siemens' current state regarding Generative AI can be described as an **exploration phase**. This is driven by the technology's novelty and Siemens' approach to the topic over the last year. **Figure 9** shows the key aspects of this exploration phase that were mentioned during the interviews and observed during the case study.

Table 3: Value Creation Areas

Functions	Business Capabilities	Opportunity Areas
Legal	Strategic & Information Management	Product / Services
Customer Service	Sales & Marketing	Front Office
Operations	Innovation & Lifecycle Management	Back office
Finance & Strategy	Supply Chain Management	Core Capabilities
Product and R&D	Manufacturing and Production	
Sales & Marketing	Service Management	
Quality Management	Project Management	
Human Resources	Enterprise Services	
Supply Chain Management		

**Figure 9:** Key aspects of the exploration phase

A consistent pattern that emerged across interviews was **creating space for innovation**. Interview partners frequently noted that they were given the resources and time necessary to explore this new technology in depth. While this seems logical for technology departments previously involved in NLP development, many business areas were also granted the resources to participate in the early ideation stage. Remarkably, many participants, particularly domain experts, highlighted that their exploration of Generative AI was driven by their own intrinsic interest. Based on this personal motivation, they were given the freedom to explore the topic within a company context (IP-13, Domain Expert).

The second theme contributing to the exploration phase is **knowledge creation**. Findings suggest that many use cases are far from deployment or are already obsolete due to the field's rapid technological advances over the last year. However, as emphasized by many interview partners, the knowledge gained during their development is a valuable asset for the future. Additionally, several participants outlined the need to invest in the area of Generative AI to stay competitive. As IP-15 puts it:

[...] we need to start building up resources as quickly as possible, because once the topic takes off, even with competitors, it will be all the more difficult to recruit new people (IP-15, CIO).

Lastly, two more themes emerged in several interviews. **Understanding the current setup** summarizes feedback that companies should carefully analyze their current pain points to apply new technologies such as Generative AI effectively. Participants pointed out that this analysis should consider both the quality of existing data and potential bottlenecks (IP-5, IT Professional), as well as the business processes where improvement offers the greatest added value (IP-10, Program Manager). Finally, **acknowledging limitations** highlights the fact that the introduction and exploration of new technologies comes with various uncertainties. One uncertainty is the rapid development of Generative AI itself, making it difficult to commit and invest in a specific solution, as it could become obsolete in a few months. Technology-inherent limitations such as hallucinations, the limited explainability of its outcomes as well as regulatory and legal considerations were named as additional constraining factors (IP-16, Head of Technology).

Besides the outlined area of exploration, interview partners expressed various requirements and challenges which can be summarized under the term **scaling**. It became apparent that for large corporations, the path from individual pilots to widely accepted approaches and processes is a critical hurdle that must be overcome to achieve sustainable added value from new technologies like Generative AI. To facilitate a better understanding of the various aspects, the results are categorized into the following dimensions:

- People
- Processes
- Technology
- Data
- Organization
- Strategy
- Communication
- Timing

While slight overlaps are possible between these dimensions (such as People and Organization or Timing and Strategy), there are distinct differences. This is also evidenced by the fact that over 500 code segments from the interviews could be assigned to the area of scaling. The main findings from the analysis are therefore presented below.

The **people** dimension deals with the influence of Generative AI on employees within the company. A dominant theme that emerged repeatedly was the belief that Generative AI will have a major impact, requiring reskilling and upskilling of the workforce (IP-1, Program Manager). Participants noted challenges stemming from the technology's rapid pace of evolution. While some colleagues adapt quickly, interviewees expressed concern that others, often from the older generations, might struggle to keep up. This could foster a negative attitude towards Generative AI, further fueled by fears of job displacement (IP-2, Program Manager).

Another hurdle identified was the lack of experts with both business knowledge and deep understanding of the technology, making value-adding applications difficult to identify (IP-3, Product Owner). Additionally, due to widespread media coverage and heightened expectations, maintaining employee motivation as initial ideas turn into concrete use cases becomes challenging (IP-5, IT Professional).

Several key factors for overcoming these hurdles were identified. Primarily, awareness must be fostered through training, workshops, or coaching, accompanied by clear expectation management. One interview partner emphasized that the psychological impact, and thus the human factor, cannot be underestimated (IP-8, CFO). Cultural differences must be considered to prevent negative consequences from the introduction of Generative AI. A simple but illustrative example is the AI-supported recording and transcription function in Microsoft Teams, which summarizes meeting content, analyzes it, and defines action items. This could lead to employees feeling uncomfortable and contributing less to future meetings.

Finally, talent management is crucial for successful scaling within the organization. Identifying key internal resources is important to recognize gaps and, if necessary, fill them with external candidates. In this context, increased cooperation with universities was mentioned as a way to attract young talent to the company (IP-15, CIO).

In terms of **processes**, the data analysis revealed that the introduction of Generative AI creates challenges similar to those faced with other technologies. This can be attributed to the complex and heterogeneous process landscape within companies like Siemens. The diverse business areas, specific requirements, and sometimes historically grown processes limit the company-wide scaling of new technologies. Interviews indicated that participants view more uniform processes and standardized approaches as major levers for scaling (IP-14, Program Manager). In this context, particular importance was placed to new software solutions whose setup would offer the possibility to consider the requirements of Generative AI at an early stage (IP-5, IT Professional).

In terms of the **technology** itself, several interesting in-

sights emerged. It became clear that many use cases failed to meet initial expectations, primarily due to the technology's current maturity and limitations for complex business processes (IP-1, Program Manager). A key aspect is the previously mentioned issue of hallucinations. Unlike previous models that provided deterministic answers, Generative AI responses are probabilistic (Radford et al., 2019, p. 2; OpenAI, 2023). It generates outputs by calculating probabilities based on its training data. The technology's maturity also presents a significant organizational hurdle. Companies must identify and evaluate different providers and solutions before approving them for internal use. This process requires time and resources and can become a bottleneck due to the rapid pace of the technological development. The data revealed several elements to address these challenges: First, Companies should strive to remain vendor and model independent. Open AI, Microsoft, Google, AWS, and open-source providers constantly release newer and better Generative AI solutions. As mentioned by IP-6, internally developed platforms should be as independent as possible to benefit from these enhancements. Second, technological stacks should be shared within the company to avoid duplicate approaches and reduce initial development efforts (IP-3, Product Owner). Third, companies must carefully weigh whether internal development is necessary or if solutions can be purchased externally. This includes cases like in-house training of large language models or use cases that might be covered by future Microsoft Copilot functionalities (IP 16, Head of Technology).

The **data** dimension plays a significant role in the successful implementation and scaling of Generative AI. This stems from the expectation that companies can achieve a competitive advantage through leveraging their own data. However, many participants expressed the need to structure and connect data within the company for meaningful use (IP-5, IT Professional). Additionally, insufficient data quality in certain areas poses a challenge. The findings suggest that a clear data strategy addressing the aforementioned aspects is necessary for successful scaling of the technology. As one interview partner mentioned, it could be beneficial to appoint a Chief Data Manager to emphasize the importance of the topic and guide it in the right direction (IP-8, CFO). The Siemens Data Cloud, an existing approach to data sharing and streamlining data processing, was cited as a positive example (IP-1, Program Manager). Overall, all participants emphasized the importance of this dimension.

Considering the complexity of large companies, the **organizational impact** of groundbreaking technologies like Generative AI can create significant challenges. This is the case for Siemens, as evidenced by the data analysis. A recurring challenge frequently addressed by the interview partners is the potential for a silo mentality, where multiple teams might work on related topics independently (IP-22, Program Manager). While decentralized structures offer individual business units freedom, they can hinder the efficient introduction of new technologies. Therefore, clear structures, a common strategy, and the sharing of technological approaches

are considered essential for successfully scaling Generative AI use cases (IP-4, Procurement).

During the exploration phase mentioned at the beginning, the aim is to build up knowledge and try out all facets of the new technology. However, converting these findings into added value for the company requires a clear **strategy**. Findings suggest this is a challenge regarding Generative AI, as there is a lack of clear metrics to measure the success of individual use cases. Especially regarding productivity-enhancing use cases, evaluation in absolute figures becomes difficult as pointed out by one interview partner (IP-10, Program Manager). Findings suggest that when developing Generative AI solutions, prioritizing business areas where the technology has the greatest potential impact seems beneficial. This aligns with recommendations from interview partners. For example, the engineering area at Siemens Mobility was identified as a potential focus due to its large share in project execution and lack of skilled railway engineers. Productivity improvements in this area could therefore yield substantial improvement potential (IP-2, Program Manager). This aligns with a key finding when analyzing individual use cases. Generative AI offers a myriad of possible applications in a company like Siemens. Consequently, it was emphasized that a considered approach to resource allocation, including deliberate decisions on individual use cases, is essential for optimal value (IP-8, CFO).

The vast attention surrounding Generative AI suggests the need for clear internal **communication** within the company. Interview partners pointed out that widespread awareness is key for effective cross-departmental use of the technology. An interesting pattern emerged during the case study: In the initial months, the desire for exchange and communication solidified into cross-departmental communities. These communities used internal platforms like Microsoft Viva Engage (**Appendix 4**) and Teams (**Appendix 5**), hosting regular information events. Existing AI structures were expanded, and their communication efforts were strengthened to meet the immense thirst for knowledge among employees and managers. One example is the “AI Attack” community, an IT group focused on streamlining the development and operation of AI solutions. After ChatGPT’s release, a separate, secure internal Generative AI Platform was established in partnership with Microsoft (see **Figure 10**).

This platform enables employees to use the technology without fear of exposing confidential data and was accompanied by an information campaign, significantly contributing to the company-wide distribution. Other communities formed in areas such as human resources, further contributing to knowledge transfer and idea exchange.

In recent months, a new trend has been observed that fits in with the previous findings, namely the establishment of cross-sector initiatives. These reflect the need and desire to bundle use cases, create more transparency and drive the topic forward strategically. Due to the size of the company, these were not only set up at corporate level but also in the individual business units in order to take account of individual needs and priorities. One pattern, expressed by almost all

interview participants, was the desire to create transparency around individual use cases. This reflects a pre-existing need within the company, as efforts to gather and showcase AI use cases were already underway. The emergence of Generative AI further amplified this need, leading to the creation of new use case collections across different areas of the company. One example at the corporate level is the so-called “Innovation Radar”, as shown in **Figure 11**. The radar shows use cases from all areas of the company and stages of development, ranging from initial ideas to lighthouse projects. The interactive display allows the user to filter by different views, such as functions, objectives, or application clusters. As part of the case study, this overview was used alongside other area-specific use case collections to verify interview results and gain a better understanding of use case examples.

Finally, a theme that can also be seen as a subtopic of the other dimensions is **timing**. This dimension encompasses all observations regarding the appropriate sequencing and approach for the introduction of Generative AI. Interview data revealed diverse perspectives on this point, which can be divided into two areas. While the company often employs a well-defined structure and milestones for piloting and introducing individual projects, such as the Siemens Industrial Copilot, it faces challenges in implementing the technology at scale. As mentioned previously, the complex systems and processes were frequently cited as reasons for this. With regard to internal productivity increases, interviewees pointed out that a few impactful use cases could generate widespread employee enthusiasm for the technology. This aligns with the author’s observations: while some employees are proactively integrating Generative AI into their daily work, the majority are not yet actively using it (**Appendix 6**).

5.6. Additional Observations

Further insights can be extracted from the conversations and data collected. A major topic area revolves around the risks associated with Generative AI and possible mitigation strategies. The Siemens legal and compliance department issued guidelines at an early stage regarding the use of external services such as ChatGPT. These included rules relating to sensitive information, illegal or unethical use, the protection of property rights, data privacy, and export control. Potential discrimination, biases, reputational risks, and questions about accountability for decisions made by Generative AI. Additionally, some participants expressed concerns about potential negative effects from anticipated increases in productivity. While many point to the skilled worker shortage and the need to streamline routine processes, others foresee more far-reaching consequences. As IP-1 puts it:

“I like the efficiency gain of it, but I also dislike the efficiency gain of it (IP-1, Program Manager)”

Various mitigation measures were mentioned to minimize risks. The most frequent was “keeping a human in the loop”, meaning that critical decisions should continue to be made or overseen by humans. Using Generative AI as a co-pilot

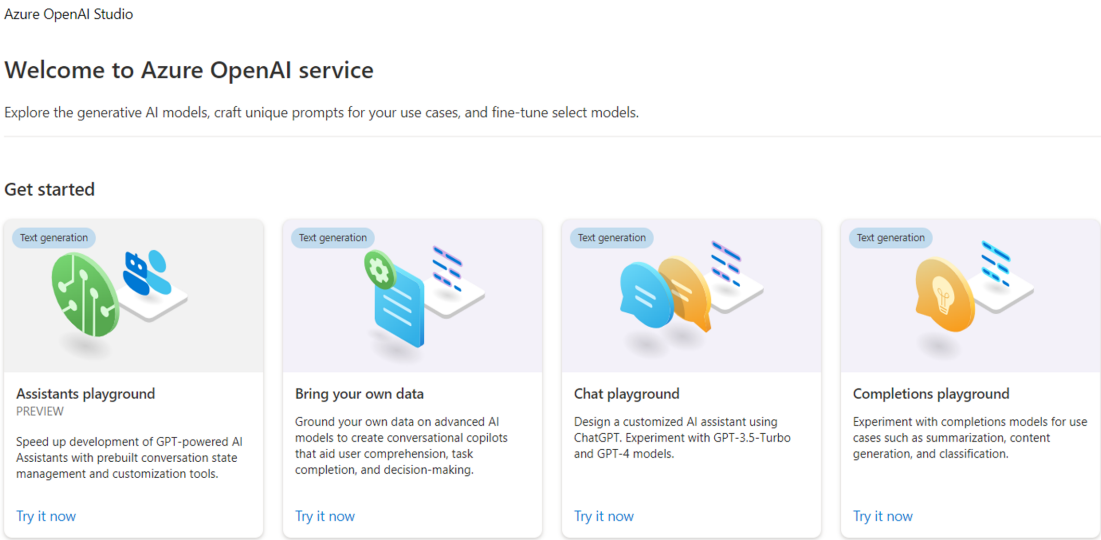


Figure 10: Siemens' internal Generative AI platform




Figure 11: Siemens' Innovation Radar for Generative AI

(though this term has been adopted by Microsoft and others for their own solutions) was frequently mentioned in this context (IP-16, Head of Technology). From the interviewees' perspective, the use of synthetic data can reduce the risk of copyright lawsuits. And while the EU Artificial Intelligence Act is seen as a further challenge, it also provides guidance for future developments in the area of Generative AI.

Finally, two additional aspects from the interviews should be mentioned. Firstly, interviewees agree that Siemens can be seen as a pioneer in implementing Generative AI within its industry. This is attributed to the early provision of infrastructure (such as the internal Generative AI platform), which allowed employees from all business units to develop

use cases. The associated knowledge building is considered valuable in initiating discussions and developing the framework for further implementations. Additionally, top management quickly embraced the topic, leading to regular discussions at internal events. External presentations, such as the Industrial Copilot at the Hannover Messe or Roland Busch's CES keynote speech, underline the importance of the technology from a corporate perspective. Lastly, another insight emerged: many interview partners see a kind of hype around Generative AI. This is particularly due to the low barrier to entry, though implementing value-adding applications for the company remains challenging. As IP-5 puts it:

“I think we all still have some way to go. On the one hand, learning how it works, but also a certain tolerance for frustration. [...] that will definitely come when the initial hype has died down and you realize that it's a bit more complicated to generate all this stuff. (IP-5, IT Professional)”

6. Discussion

This case study explores the value creation potential of Generative AI within a large, multinational corporation. Spanning a one-year observation period and drawing on 23 in-depth interviews with experts and managers across various business areas, it offers a comprehensive view of the processes and challenges associated with introducing this transformative technology. The following section will delve into the key findings and their implications for both researchers and practitioners. This analysis will address the central research questions and highlight the study's theoretical and practical contributions. The chapter will conclude with a critical reflection, acknowledging the study's limitations and outlining potential avenues for future research.

6.1. Technological Innovation and Usability

When OpenAI launched ChatGPT in November 2022, the public was able to experience the advances in the field of artificial intelligence for the first time. Based on the outlined work by Brown et al. (2020), Radford and Narasimhan (2018), and Vaswani et al. (2017) and many others, ChatGPT showed what was possible with the help of “artificial intelligence”. Some studies pointed to the possible significant influence of LLMs on different occupational groups (for example, Eloundou et al., 2023, p. 11) while others recognized the ability of Generative AI to take on creative tasks that previously could only be performed by humans (Feuerriegel et al., 2023, p. 1). Exploring **Research Question 1** on how Generative AI differs from previous models in terms of architecture and functionality, two key findings emerge from a business perspective. Firstly, the technological innovation enables companies like Siemens to explore new application areas and thus opportunities to generate additional business value. This is particularly important for existing products and software solutions that can be improved with the help of Generative AI. Examples that have been outlined in the case study are the Siemens Industrial Copilot integrated in Siemens' automation environment or Siemens' low-coding platform Mendix that incorporates Generative AI capabilities. While these are customer-facing products, the adaptability of Generative AI appears to be useful for internal applications as well. Use cases can be realized with significantly less effort, i.e., with a lower investment volume and in less time compared to previous models, thus enabling more business areas to deploy AI solutions to streamline their processes. Although long-term productivity improvements are not yet demonstrable, the case study has shown that efficiency gains are anticipated across all business. This starts

with information retrieval systems which reduce the time spent for searching specific information and continues with assistant systems which actively support the user and help improve quality in predefined tasks. From the company's point of view, external solutions in particular can offer added value alongside in-house developments. Building on partnerships with software providers, such as Microsoft, companies like Siemens can benefit directly from new AI functionalities. This aligns with corporate goals, as the widespread adoption of external software solutions maximizes the reach of potential improvements. In the case of Microsoft, Siemens was able to evaluate the Microsoft Copilot at an early stage and make it available to selected employees as part of an internal pilot. This allowed the company to assess the Generative AI functionalities and build knowledge within the organization, laying the foundation for future implementation. Addressing the central research question, another key finding is the notable usability of Generative AI. Its ability to execute tasks based on natural language prompts offers a significantly streamlined user experience for interacting with AI systems. This, coupled with the accessibility of platforms like OpenAI's ChatGPT (and subsequent internal solutions), expands the scope of potential applications. For large corporations, this presents a significant shift, empowering employees without specialized AI knowledge to directly utilize the technology.

6.2. Emerging Patterns of Value Creation

Based on the case study results, the value creation potential of Generative AI has been outlined in Chapter 5. Use cases were clustered into broader schemes and key drivers were identified. Moreover, application areas were discovered which help companies like Siemens to structure and prioritize its use cases. With this, the case study supports closing the gap between scientific concepts and real-world application. It provides an answer to **Research Question 2**, namely in which ways Generative AI presents new avenues for value creation. Based on the technological innovation and usability as described above, the clustering of use cases in particular helps to understand the potential added value. Contrasted with the earlier work on value creation mechanisms, such as the process model of ML value creation by Shollo et al. (2022) or the four sources of value creation by Borges et al. (2021), similarities but also divergences become visible (see **Table 4**). Most noticeable, the use of AI for information retrieval can be seen as a new area for value creation which was not covered before. This can be traced back to the technological advances of Generative AI and its capability to process and provide answers from large data sets.

Another segment which was not touched on before, is coding support. While it can be argued that this could be seen as a subset of the cluster assistants, its far-reaching implications can serve as a reason for presenting it separately. With their knowledge of all major programming languages, Generative AI models can help to generate, adjust, test, and optimize code automatically, changing the way programmers deal with tasks and business problems. These unique capabilities of Generative AI could not be observed by researchers in

Table 4: Comparison of case study results with previous research

Case Study Clusters	Shollo et al. (2022): Shifting ML value creation mechanisms	Borges et al. (2021): The strategic use of artificial intelligence in the digital era
Assistants - Information Retrieval	N/A	N/A
Data Synthesis	Knowledge Creation	Decision Support
Assistants - Coaches & Helpers	Task Augmentation	Customer and Employee Engagement
Autonomous Assistants	Autonomous Agents	Automation
Content Creation	N/A	New Products and Services
Coding Support	N/A	N/A

2021 and 2022 respectively and therefore represent a valuable supplement to existing research.

6.3. Strategic and Operational Deployment

Besides the outlined avenues for value creation, supporting strategies and considerations have been identified. To address **Research Question 3**, which explores ways to harness Generative AI’s full value creation potential, the following section discusses these findings in detail. Based on the insights gained regarding exploration and scaling, a new framework is proposed (**Figure 12**). It illustrates the connection between these phases, along with the necessary tasks to effectively exploit Generative AI’s value creation potential. The hourglass shape symbolizes the critical interface between the two phases. It emphasizes the need for a business impact assessment to successfully scale use cases across the organization. This assessment requires weighing the costs and benefits for each use case individually. Since companies seek the best possible use of their limited resources, not all ideas from the exploration phase will necessarily be implemented. In order to determine the added value, subsequent questions should be addressed, namely who coordinates the implementation, how much development effort is required, which accompanying change management measures are necessary and who is responsible for the care, maintenance, and costs of the new solution.

The business impact assessment goes hand-in-hand with strategic considerations. For large companies, it is essential to view AI initiatives within a broader context and align them with corporate strategy (Kitsios & Kamariotou, 2021, p. 6). This is especially important since the added value of Generative AI use cases might not be immediately measurable. Companies should consciously define core areas to steer resources effectively. In this context, Generative AI presents a new challenge for managers and strategists. While past AI initiatives were more manageable and clearly definable, technological innovation now offers a myriad of potential value creation opportunities. Evaluating and filtering these, along with defining a clear strategy, becomes an essential task for successfully utilizing Generative AI. The framework suggests

main tasks for each phase. In the exploration phase, these include creating awareness, sharing knowledge, building communities, and providing a common technology stack. To support the business impact assessment and strategic alignment, key tasks involve creating transparency, minimizing macro-risks, and allocating resources to identified use cases. Finally, to successfully scale solutions, companies should focus on harmonizing their data landscape, adjusting processes, and offering training and guidance to employees.

6.4. Theoretical Contributions

The results of this case study contribute to existing research in the fields of value creation, AI adoption, and capabilities. This thesis offers a valuable contribution by analyzing proposed frameworks by Borges et al. (2021) and Shollo et al. (2022) within the specific context of Generative AI. The analysis identifies areas where these models could be further expanded. Additionally, the framework by Uren and Edwards (2023), with its four lenses People, Processes, Technology, and Data proved highly applicable. This case study expands upon their model by detailing specific considerations within each category relevant to Generative AI implementation. Furthermore, new dimensions emerged for analysis, including the importance of timing, organizational setup, strategy, business, value and communication. The case study aligns with the findings of Caner and Bhatti (2020), highlighting the multiple perspectives from which AI can be viewed. For example, the results suggest that while Generative AI has overcome certain limitations like data labeling or generalizability of learning, others, such as biases and non-explainability, remain a challenge. Another interesting avenue for future research could be to investigate further application areas of Generative AI, as its primary use is currently in augmentation tasks, with autonomous systems seeming less prevalent. In conclusion, this thesis adds to the existing body of knowledge by refining and extending established models in the context of Generative AI. This detailed exploration offers valuable insights that can advance theoretical understanding and guide the development of more comprehensive frameworks for AI adoption.

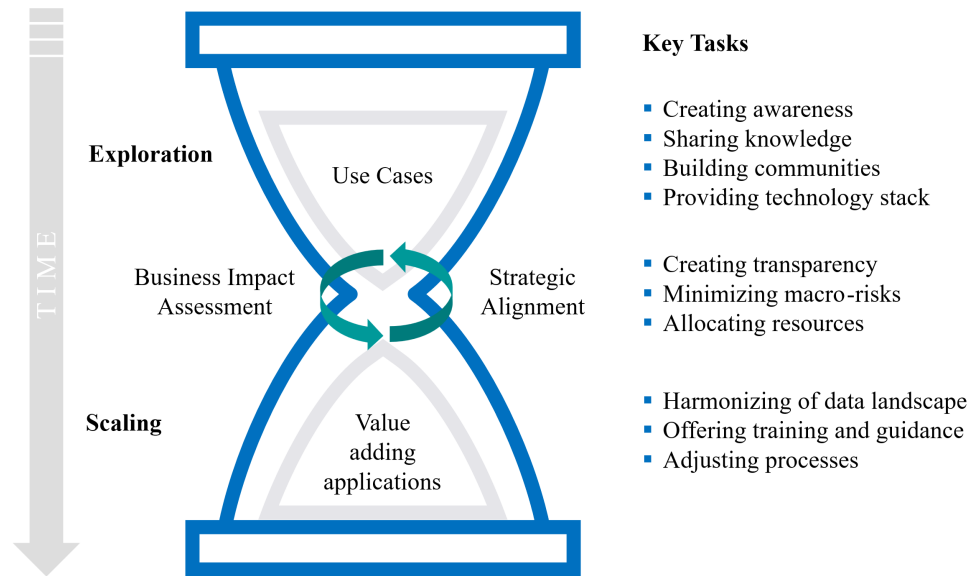


Figure 12: Proposed framework for exploration and scaling of Generative AI

6.5. Practical Implications

Merging the results from the case study and theoretical concepts on AI value creation and capabilities, practical implications emerge. In the following section, these should be outlined in more detail, providing companies with a clearer understanding of Generative AI and the required actions to create added value.

Firstly, companies should **create room for innovation**. Generative AI is unique in that its core functionalities (i.e., summarization, translation, and content creation) can be used by almost any employee. This empowers more people, even those without a background in IT or data science, to develop business ideas within their respective domains. Companies should embrace this opportunity and actively support the ideation and exploration of new Generative AI solutions.

An important prerequisite for this is the **provision of the required infrastructure**. With regard to Generative AI, this comprises company-internal platforms and software packages that employees can use safely and securely, without the risk of exposing sensitive company information externally. At the same time, it is important to highlight the limitations of the technology to avoid exaggerated expectations.

Secondly, companies should **fight for transparency**. Large corporations with their decentralized structures tend to develop similar solutions in multiple business areas. With Generative AI, this challenge becomes even more apparent as use cases can be developed within days or weeks. While exploration should be encouraged, creating transparency over all initiatives is key for companies; otherwise, resources risk being wasted on redundant solutions. Employees and departments should therefore be encouraged to share their knowledge and solutions with a broader audience. This should not be taken lightly: active management of potential use cases and clear guidance for employees are crucial to

reduce isolated solutions, frictional losses, and unfulfilled expectations. If possible, top management should enforce the needed structures to provide a company-wide framework.

Thirdly, a **clear AI strategy** is needed to steer the company's resources in the right direction. As outlined before, Generative AI offers multifold opportunities for integration into processes and enhancement of existing products. However, companies should carefully evaluate each use case and its contribution to the overall company strategy. Regarding customer-facing solutions, the inherent technological risks need to be analyzed in detail. For internal processes, the added value of Generative AI must be weighed against increased complexity and consequential costs. To guide teams effectively, companies should provide guidance for their middle management, enabling them to evaluate new ideas within a broader context.

Finally, companies should focus on **bridging the gap between Generative AI knowledge and operational business units**. Since Generative AI use cases can increasingly be driven directly by these units without extensive IT or data science involvement, it is wise to empower them. Companies should consider strengthening existing departments (such as innovation, operational, or business excellence) or establishing dedicated Generative AI resources across the organization to maximize the technology's potential.

6.6. Study Limitations

Although providing valuable insights into the value creation opportunities of Generative AI, this case study is not without its limitations. The first limitation lies in the design of the study. As a single case study, it offers unique and comprehensive insights into the object of investigation; however, its findings cannot be broadly generalized. The second limitation is rooted in the author's active involvement in the

exploration of Generative AI in the company. While this allowed access to information that would normally not have been accessible (for example, the participation in external and internal conferences or the regular exchange with data scientists on technical aspects of Generative AI), biases and personal perceptions could impact the final results, despite the author's best efforts to avoid unintended interferences. The third limitation concerns the data used for the case study. Due to the novelty of Generative AI, existing concepts of AI value creation, adoption, and capabilities were used, concepts that might not fully reflect the latest technological advances. The selection of concepts was based on a structured approach; however, it cannot be ruled out that other unconsidered scientific work might have provided additional value. Regarding the data collected in the company, it should be noted that the selection of interview partners might have influenced the results. An attempt was made to cover as many areas and functions as possible, but a selection was necessary due to the substantial number of potential contacts.

In conclusion, this work represents a snapshot in time. The field of Generative AI is evolving rapidly. While the focus at the beginning of the case study was on textual understanding, new models now generate both images and videos. Despite the described limitations, this work aims to contribute to ongoing research on the value creation potential of Generative AI.

6.7. Future Research Opportunities

Building on the outlined limitations, future research could verify and expand the findings of this case study. Firstly, another case study in a similar setup could verify and strengthen the results, potentially adding to the proposed framework for exploration and scaling. Additionally, quantitative data would help deepen the understanding of a company's requirements and anticipated challenges when introducing Generative AI. Surveys among managers and employees could provide valuable insights into critical success factors. Moreover, more scientific work is needed to explain the unique capabilities of Generative AI compared to earlier machine learning models, and to integrate these findings into existing conceptual frameworks. Finally, future research should investigate the impact and value creation potential of the latest technological advances. With new models such as Google's Gemini, providing a context window of up to one million tokens, recently developed solutions for information retrieval could become obsolete. This and other advances such as OpenAI's text-to-video model Sora continually open new avenues for research in this field.

7. Conclusion

In this study, the value creation opportunities of Generative AI were explored. Based on a single case study design, Siemens was chosen as the object of investigation as it offered a unique opportunity to study the phenomenon in depth. Over a period of one year, a considerable amount

of time was spent observing and participating in numerous events to achieve a better understanding of the technology and Siemens' approach to introducing and exploiting its value creation potential. This was supplemented by 23 comprehensive interviews with key stakeholders, providing a complete view of use cases, activities, perspectives, and considerations.

The findings suggest that Generative AI presents new value creation opportunities, driven by technological advances and usability that bring AI closer to domain experts. Use cases have been identified across various areas and clustered into four main groups. Smart assistants with various levels of complexity form the majority of observed pilots, along with lighthouse projects in different business areas. While many applications offer potential added value, the study acknowledges that a considerable number of projects are still in their early stages.

Contributing to the ongoing discourse in AI research on value creation and required capabilities, different frameworks were applied to explain the case study findings. For example, addressing the capabilities for the successful implementation of AI, the outlined dimensions from previous research reappeared in the data. Other concepts could not be fully confirmed by the findings, opening room for further research. For example, existing work on value creation mechanisms was not able to explain all value creation clusters uncovered in the case study.

A key contribution of the thesis can be seen in its long-term observation of a company in the early adoption phase of a new technology. Starting in the early discovery stage only a few months after the release of Generative AI to the public, Siemens' steps to introduce, manage, and profit from the technology could be observed over a year. Building on the gathered knowledge, a new framework is proposed to explain the actions and strategies revealed in the case study. It introduces an exploration phase as an important time period in the adoption of Generative AI. Moreover, it identifies essential tasks to capitalize on the opportunities presented in this phase and outlines the necessary steps for a successful transition from exploration to scaling.

To conclude, the thesis provides practical implications and highlights potential areas for future research. Acknowledging its limitations, the thesis offers unique insights into a company's approach to Generative AI and contributes to the ongoing discourse on the potential value creation opportunities of the technology.

References

- Akkiraju, R., Sinha, V., Xu, A., Mahmud, J., Gundecha, P., Liu, Z., Liu, X., & Schumacher, J. (2018). Characterizing machine learning process: A maturity framework. <https://doi.org/10.48550/arXiv.1811.04871>
- Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural Machine Translation by Jointly Learning to Align and Translate. <https://doi.org/10.48550/arXiv.1409.0473>
- Bastian, M. (2023). GPT-4 has more than a trillion parameters. Retrieved March 25, 2023, from <https://the-decoder.com/gpt-4-has-a-trillion-parameters/>

- Blanka, C., Krumay, B., & Rueckel, D. (2022). The interplay of digital transformation and employee competency: A design science approach. *Technological Forecasting and Social Change*, 178, 121575. <https://doi.org/10.1016/j.techfore.2022.121575>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., & Liang, P. (2021). On the Opportunities and Risks of Foundation Models. <https://doi.org/10.48550/arXiv.2108.07258>
- 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, 1–16. <http://doi.org/10.1016/j.ijinfomgt.2020.102225>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- Brock, J. K.-U., & Wangenheim, F. v. (2019). Demystifying AI: What Digital Transformation Leaders Can Teach You about Realistic Artificial Intelligence. *California Management Review*, 61(4), 110–134. <https://doi.org/10.1177/1536504219865226>
- Brown, T. B., Mann, B., Ryder, N., 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., Winter, C., & Amodei, D. (2020). Language Models are Few-Shot Learners. <https://doi.org/10.48550/arXiv.2005.14165>
- Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative AI at Work. <https://doi.org/10.3386/w31161>
- Caner, S., & Bhatti, F. (2020). A Conceptual Framework on Defining Businesses Strategy for Artificial Intelligence. *Contemporary Management Research*, 16(3), 175–206. <https://doi.org/10.7903/cmr.19970>
- Davenport, T. H., & Ronanki, R. (2018). Artificial Intelligence for the Real World: Don't start with moon shots.
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. <https://doi.org/10.2139/ssrn.4573321>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., & Wright, R. (2023). Opinion Paper: "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. <http://arxiv.org/pdf/2303.10130v4>
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2023). Generative AI. *Business & Information Systems Engineering*. <https://doi.org/10.1007/s12599-023-00834-7>
- Gartner. (2023). What Is Your Enterprise's AI Ambition? <https://www.gartner.com/en/documents/4836831>
- Glaser, B. G., & Strauss, A. L. (1967). *The Discovery of Grounded Theory*. Routledge. <https://doi.org/10.4324/9780203793206>
- Goldman Sachs. (2023). Generative AI could raise global GDP by 7%. <https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>
- Graves, A. (2012). *Supervised Sequence Labelling with Recurrent Neural Networks* (Vol. 385). Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-24797-2>
- Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for boltzmann machines. *Cognitive Science*, 9(1), 147–169. [https://doi.org/10.1016/S0364-0213\(85\)80012-4](https://doi.org/10.1016/S0364-0213(85)80012-4)
- Hochreiter, S., Bengio, Y., Frasconi, P., & Schmidhuber, J. (2001). Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. *A field guide to dynamical recurrent neural networks*. <https://doi.org/10.1109/ICASSP.2001.940859>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Kim, Y., Denton, C., Hoang, L., & Rush, A. M. (2017). Structured Attention Networks. <https://doi.org/10.48550/arXiv.1702.00887>
- Kitsios, F., & Kamariotou, M. (2021). Artificial Intelligence and Business Strategy towards Digital Transformation: A Research Agenda. *Sustainability*, 13(4), 2025. <https://doi.org/10.3390/su13042025>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Lee, J., Suh, T., Roy, D., & Baucus, M. (2019). Emerging Technology and Business Model Innovation: The Case of Artificial Intelligence. *Journal of Open Innovation: Technology, Market, and Complexity*, 5(3), 44. <https://doi.org/10.3390/joitmc5030044>
- Lichtenthaler, U. (2019). An Intelligence-Based View of Firm Performance: Profiting from Artificial Intelligence. *Journal of Innovation Management*, 7(1), 7–20. https://doi.org/10.24840/2183-0606_007.001_0002
- Luger, G. F., & Stubblefield, W. A. (1998). *Symbolic Computation* (3rd). Addison-Wesley.
- McCarthy, J. (1958). Programs with common sense: Proceedings of the Symposium on Mechanisation. London 77–84.
- McCarthy, J., & Feigenbaum, E. A. (1990). In Memoriam: Arthur Samuel: Pioneer in Machine Learning. *AI Magazine*, 11(3), 10. <https://doi.org/10.1609/aimag.v11i3.840>
- Mitchell, T. M. (1997). *Machine Learning*. McGraw-Hill.
- Myers, M. D. (2020). *Qualitative Research in Business & Management* (Third). SAGE Publications Ltd.
- Ng, A. Y., & Jordan, M. I. (2001). On Discriminative vs. Generative Classifiers: A comparison of logistic regression and naive Bayes.
- Nilsson, N. J. (2013). *The Quest for Artificial Intelligence*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511819346>
- OpenAI. (2015). Introducing OpenAI. <https://openai.com/blog/introducing-openai>
- OpenAI. (2022). Introducing ChatGPT. <https://openai.com/blog/chatgpt>
- OpenAI. (2023). GPT-4 Technical Report. <http://arxiv.org/pdf/2303.08774v3>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. <http://arxiv.org/pdf/2203.02155v1>
- Peng, B., Alcaide, E., Anthony, Q., Albalak, A., Arcadinho, S., Biderman, S., Cao, H., Cheng, X., Chung, M., Grella, M., GV, K. K., He, X., Hou, H., Lin, J., Kazienko, P., Koon, J., Kong, J., Koptyra, B., Lau, H., & Zhu, R.-J. (2023). RWKV: Reinventing RNNs for the Transformer Era. <https://doi.org/10.48550/arXiv.2305.13048>
- Radford, A., & Narasimhan, K. (2018). Improving Language Understanding by Generative Pre-Training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language Models are Unsupervised Multitask Learners. <https://arxiv.org/pdf/1909.01317v2.pdf>
- Rao, A. (2017). A Strategist's Guide to Artificial Intelligence. <https://www.strategy-business.com/article/A-Strategists-Guide-to-Artificial-Intelligence>
- Reilly, A., Depa, J., & Douglass, G. (2019). AI: Build to Scale: From Experimental to Exponential. <https://www.accenture.com/content/dam/accelenture/final/a-com-migration/thought-leadership-assets/accelenture-built-to-scale-pdf-report.pdf>
- Reim, W., Åström, J., & Eriksson, O. (2020). Implementation of Artificial Intelligence (AI): A Roadmap for Business Model Innovation. *AI*, 1(2), 180–191. <https://doi.org/10.3390/ai1020011>
- Rubin, H. J., & Rubin, I. S. (2012). *Qualitative interviewing: The art of hearing data*. SAGE Publications.

- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536. <https://doi.org/10.1038/323533a0>
- Russell, S. J., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach* (Fourth). Pearson.
- Samuel, A. L. (1959). Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, 3(3), 535–554. <https://doi.org/10.1147/rd.33.0210>
- Shollo, A., Hopf, K., Thiess, T., & Müller, O. (2022). Shifting ML value creation mechanisms: A process model of ML value creation. *The Journal of Strategic Information Systems*, 31(3), 1–20. <https://doi.org/10.1016/j.jsis.2022.101734>
- Siemens. (2023). Annual Report 2023. <https://www.siemens.com/global/en/company/investor-relations/events-publications-ad-hoc/annualreports.html>
- Siemens. (2024). Siemens - Who we are. <https://www.siemens.com>
- Simmons, A. B., & Chappell, S. G. (1988). Artificial intelligence-definition and practice. *IEEE Journal of Oceanic Engineering*, 13(2), 14–42. <https://doi.org/10.1109/48.551>
- Sutton, R. S., & Barto, A. (2018). *Reinforcement learning: An introduction* (Second edition). The MIT Press.
- Trunk, A., Birkel, H., & Hartmann, E. (2020). On the current state of combining human and artificial intelligence for strategic organizational decision making. *Business Research*, 13(3), 875–919. <https://doi.org/10.1007/s40685-020-00133-x>
- Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind*, 59(236), 433–460. <https://doi.org/10.1093/mind/LIX.236.433>
- Uren, V., & Edwards, J. S. (2023). Technology readiness and the organizational journey towards AI adoption: An empirical study. *International Journal of Information Management*, 68, 102588. <https://doi.org/10.1016/j.ijinfomgt.2022.102588>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need. <http://arxiv.org/pdf/1706.03762v5>
- Voita, E., Sennrich, R., & Titov, I. (2019). The Bottom-up Evolution of Representations in the Transformer: A Study with Machine Translation and Language Modeling Objectives. <https://doi.org/10.48550/arXiv.1909.01380>
- Weber, M., Engert, M., Schaffer, N., Weking, J., & Krcmar, H. (2023). Organizational Capabilities for AI Implementation—Coping with Inscrutability and Data Dependency in AI. *Information Systems Frontiers*, 25(4), 1549–1569. <https://doi.org/10.1007/s10796-022-10297-y>
- Yin, R. K. (2018). *Case study research and applications: design and methods* (6th ed.). SAGE.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., & Irving, G. (2019). Fine-Tuning Language Models from Human Preferences. <https://doi.org/10.48550/arXiv.1909.08593>