Università degli Studi di Padova

DIPARTIMENTO DI TECNICA E GESTIONE DEI SISTEMI INDUSTRIALI

(DTG)

Master thesis in Management Engineering

Application of Generative AI to the business context: analysis and comparison

Supervisors :

Prof. Luigi Salmaso

Dr. Nicolò Biasetton

Dr. Elena Barzizza

Master candidate:

Abdelhaleem Jalal Abdelhaleem Ahmed

Student ID: 2073558

Academic Year: 2024-2025

Abstract

The rapid and widespread deployment of efficient artificial systems intelligence (AI) is deemed to have significant impact on businesses as it can lead to improved efficiency, increased sales and marketing, and higher profitability. Especially, the utilization of Generative Artificial Intelligence (GenAI) technologies can affect several areas and can be disruptive in terms of competition between companies. However, their effectiveness in specific cases is still controversial. In addition the use of these systems may also rise legal, ethical, and ethical issues such as copyright infringement and the production of fake and counterfeit goods information. It is therefore important for organizations to understand the real benefits of AI and minimize possible harm. The purpose of this thesis is to present an updated illustration of GenAI systems and their potential and problems. Also, to better understand the effectiveness of some existing systems, an empirical analysis of how they function in translating business documents is conducted.

Keywords: ChatGPT, Generative Artificial Intelligence, AI productivity, GenAI technology

Preface

In the name of Allah, the Most Gracious, the Most Merciful. Alhamdulillah, all praises to Allah for the strength and blessings that enabled me to complete this thesis.

My objective in conducting this research was to examine the practical applications and paradoxes of GenAI, as well as to develop an implementation guide for businesses seeking to integrate this technology effectively. Through a thorough analysis of sectoral use cases and emerging trends, this work aims to contribute to a more structured and responsible approach to GenAI adoption.

I extend special appreciation to my supervisor, Prof. Luigi Salmaso, Dr. Nicolò Biasetton and Dr. Elena Barzizza for their unwavering guidance and support. Their invaluable feedback and suggestions throughout the thesis have greatly contributed to its success. This project has provided me with invaluable real-world experience, and I am grateful for the opportunity. I would also like to express my gratitude to UNIPD and specially the DTG department for their efforts in providing us with essential information and facilitating our journey, enabling us to effectively implement our education into real-time project design and analysis.

Contents

List of Figures:

List of Tables:

Chapter 1

Introduction

1. Introduction

1.1 Background and Context

Generative Artificial Intelligence (GenAI) is redefining the technological landscape by introducing unparalleled capabilities in content creation, automation, and decision-making processes. Unlike traditional artificial intelligence systems that focus on predictive analytics, GenAI possesses the ability to generate entirely new, contextually relevant outputs in the form of text, images, audio, video, and code. This capability is underpinned by advancements in Transformer-based architectures, Variational Autoencoders (VAEs), and Generative Adversarial Networks (GANs), allowing GenAI to enhance productivity, creativity, and efficiency across diverse sectors.

The rapid adoption of GenAI is fueled by the increasing sophistication of Large Language Models (LLMs) such as GPT-4, PaLM 2, and open-source frameworks like LLaMA. These models, trained on vast datasets, are disrupting conventional workflows in industries ranging from healthcare and finance to manufacturing, marketing, and engineering design. Businesses now leverage GenAI to optimize processes, streamline decision-making, and create innovative products with unprecedented speed and precision (Saha, 2024; Cano-Marin, 2024).

Artificial intelligence (AI) has become a prominent topic in research and practice across a wide range of disciplines. At the center of inquiry, particularly in information systems (IS) research, is the socio-technical impact of AI as a phenomenon (Ågerfalk et al., 2022; Berente et al., 2021). These technologies have revolutionized the way we use data and make decisions, opening up previously unexplored social and economic possibilities (Fügener et al., 2021; J.Li et al., 2021; van den Broek et al., 2021). But as we enter this new era of generative AI, it is more important than ever to comprehend the fundamental ideas and differences that make GenAI a quickly developing field of technology. In the following sections, a quick summary of AI and its subfields is reported, including related areas such as machine learning (ML) and deep learning (DL), in order to distinguish GenAI from other AI concepts and offer a basic conception (Fig. 1.1). Next, we go into further detail about how DL has fueled the creation of deep generative models to enable unique GenAI features and capabilities, which has eventually led to an even wider range of AI applications with great potential for different fields.

However, while GenAI promises vast economic and operational benefits, it also introduces critical challenges, including algorithmic bias, data privacy risks, workforce displacement, and regulatory concerns. The ability to harness GenAI effectively while mitigating these risks requires a multidisciplinary approach encompassing machine learning, engineering management, and business strategy.

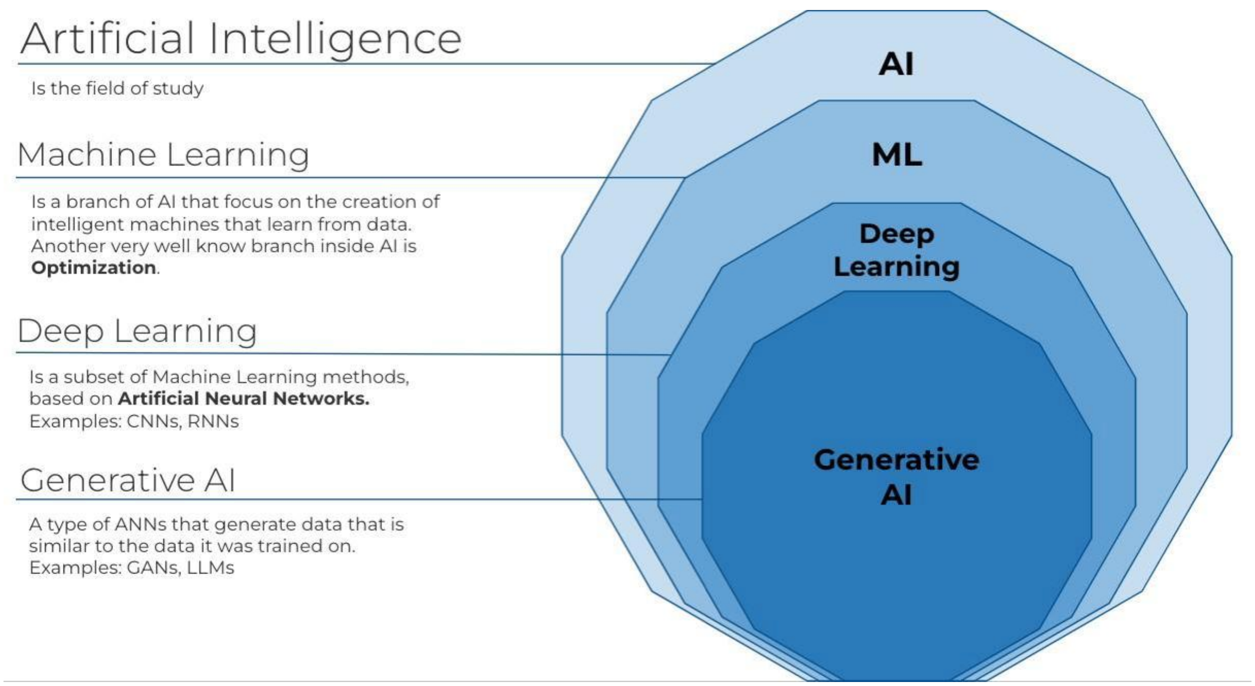

Figure 1 Generative AI and other AI concepts[1]

[1] From internet https://ghazalabioinfodr.medium.com/ai-ml-and-dl-309700212b94

Compared to the deterministic results of discriminative AI, or variance, the outputs of generative AI models are by design probabilistic and non-replicable (Weisz et al., 2023). A GenAI application will generate different outputs for the same input prompt each time, but the results remain valid and contextually relevant. However, different input prompts can still result in the same outcome. As a result, creating a relevant prompt that produces the intended result requires trial and error, such as rewording textual prompts using the same phrases. The field of prompt engineering is concerned with methodically creating prompts to enhance outputs that are produced (Liu & Chilton, 2022).

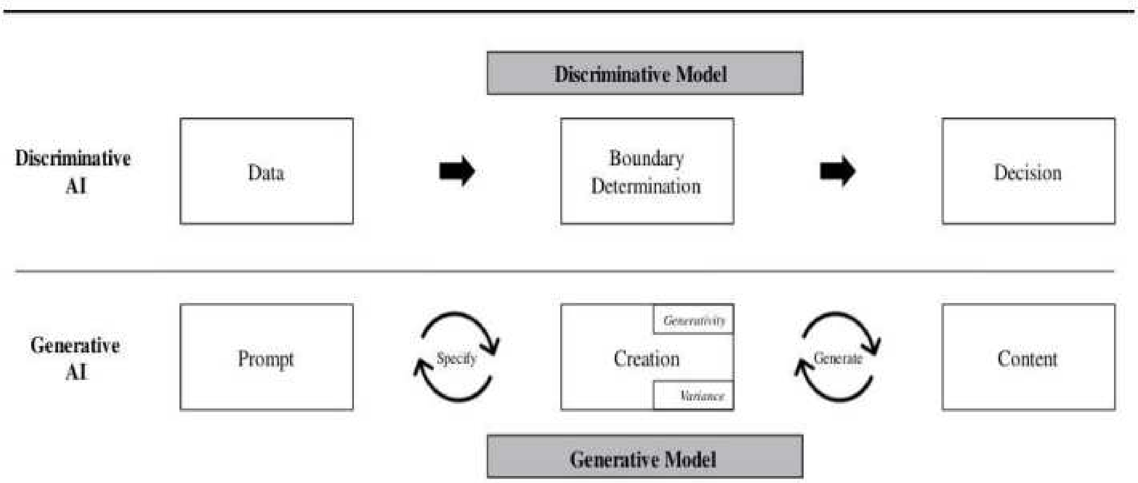

Figure 2 Procedural differences of generative AI and discriminative AI[2]

The increasing prominence of GenAI in industry necessitates a deeper exploration of its capabilities, applications, and limitations. McKinsey & Company (2023) estimates that GenAI could contribute between

[2] Procedural differences of generative AI and discriminative AI Generative AI and other AI concepts Şahin, O., & Karayel, D. (2024). Generative Artificial Intelligence in Business and Industry: A Systematic Review on the Threshold of Transformation. Journal of Smart Systems Research, 5(2), 156–175. https://doi.org/10.58769/joinssr.1597110

$2.6 trillion and $4.4 trillion annually to the global economy. This economic impact stems from its ability to:

- Enhance automation in industries such as manufacturing, customer service, and supply chain management.
- Accelerate product development by generating optimized designs, software code, and engineering solutions.
- Augment decision-making through real-time data analysis and predictive modeling.
- Enable human-AI collaboration by integrating GenAI into creative industries, research, and education.

Despite these advantages, the adoption of GenAI remains uneven across industries due to technological, ethical, and organizational challenges. While Industry 4.0 emphasized digital transformation and automation, the emerging Industry 5.0 is focused on human-centric AI, emphasizing the need for ethical governance, workforce reskilling, and sustainable AI deployment (Mohammed & Skibniewski, 2023).

Generative AI comprises various deep learning architectures, each with unique mechanisms for creating synthetic data. Among the most prominent models are Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and Transformer-based models (TBMs).

GANs operate through a competitive process between two neural networks: a generator and a discriminator. The generator creates synthetic data, while the discriminator assesses whether the generated data is real or artificial. This adversarial process continues until the generator produces data indistinguishable from real-world samples (Şahin, O., & Karayel, D. 2024). The ability of GANs to generate high-quality images and text has led to their widespread use in fields such as creative design, medical imaging, and autonomous systems.

VAEs, in contrast, rely on probabilistic modeling to encode input data into a latent space before reconstructing meaningful outputs. Unlike GANs, which operate on adversarial principles, VAEs are more structured in their approach to data generation, making them particularly useful for applications in unsupervised representation learning (Şahin, O., & Karayel, D. 2024). VAEs are extensively used in drug discovery, anomaly detection, and data compression due to their ability to handle high-dimensional datasets.

Transformer-based models (TBMs) have revolutionized natural language processing and general AI capabilities. Introduced in 2017, TBMs utilize self-attention mechanisms to process and generate sequences of data, enabling applications in text translation, speech recognition, and even DNA sequence analysis (Şahin, O., & Karayel, D. 2024). The efficiency and scalability of Transformer architectures have solidified their role as a foundation for Large Language Models (LLMs) such as GPT-4 and PaLM 2.

Another crucial development in generative modeling is diffusion models, which operate through a two-phase process: forward diffusion and reverse diffusion. Forward diffusion introduces Gaussian noise into an input dataset, gradually degrading the data. The reverse process then reconstructs the original data by learning patterns iteratively. While diffusion models produce high-quality outputs, their computational requirements remain a significant challenge (Şahin, O., & Karayel, D.2024).

1.1.1 Evolution from Rule-Based AI to Generative Models

The trajectory of AI development has transitioned from rule-based systems to generative models, marking a fundamental shift in how machines process and create data. The origins of AI can be traced back to Alan Turing's work in the early 1950s, where he proposed the notion of machine intelligence through the famous

Turing Test (Kanbach, D. K., Heiduk, L., et al., 2024). This early framework set the stage for the development of rule-based expert systems in subsequent decades.

In the 1960s, Weizenbaum introduced ELIZA, an early chatbot that mimicked human-like conversation using pre-defined scripts. ELIZA's ability to simulate human responses laid the groundwork for natural language processing (NLP) models that followed (Kanbach, D. K., Heiduk, L., et al., 2024). However, these systems were rigid and lacked the ability to generate original content beyond their programmed rules.

With advancements in computational power and data availability, AI researchers shifted towards neural networks and deep learning models in the early 2000s. The rise of machine learning algorithms—particularly in areas like computer vision, speech recognition, and fraud detection—paved the way for more dynamic and scalable AI solutions (Kanbach, D. K., Heiduk, L., et al., 2024).

The true breakthrough occurred with the introduction of Transformer-based architectures. By eliminating the constraints of traditional rule-based AI, Transformers enabled large-scale content generation, contextual awareness, and adaptation to dynamic inputs (Kanbach, D. K., Heiduk, L., et al., 2024). The release of GPT-2 in 2019 marked the onset of Generative AI's disruptive capabilities, leading to the widespread adoption of LLMs in various industries.

1.1.2 Key Enabling Technologies

The expansion of Generative AI has been facilitated by several enabling technologies, with Large Language Models (LLMs), self-attention mechanisms, and high-performance computing playing central roles.

LLMs have redefined AI's ability to process, generate, and understand human language. By training on massive datasets, these models can classify, edit, summarize, and generate text-based content with unprecedented fluency and coherence (Chui, M., Hazan, E., et al., 2023). Their applications extend beyond NLP, influencing fields such as robotics, content automation, and personalized learning systems.

A key technological advancement underlying LLMs is the self-attention mechanism in Transformer models. Unlike recurrent neural networks (RNNs), which process sequential data linearly, self-attention allows models to analyze entire sequences simultaneously, significantly improving efficiency (Chui, M., Hazan, E., et al., 2023).

However, the growing complexity of LLMs has led to increased computational demands. Training state-of-the-art generative models requires substantial GPU clusters, which can create bottlenecks in development and deployment (Chui, M., Hazan, E., et al., 2023). Addressing these challenges will require advancements in energy-efficient AI architectures and optimized training paradigms.

The evolution of Generative AI represents a paradigm shift from rule-based automation to self-learning, generative capabilities. By leveraging GANs, VAEs, Transformers, and diffusion models, businesses can unlock new levels of creativity, automation, and efficiency. However, challenges such as computational scalability and ethical concerns remain pressing issues. Understanding these technological underpinnings is essential for industry leaders seeking to harness the full potential of GenAI while navigating its complex landscape responsibly.

1.2 Research Objectives and Scope

This study aims to analyze the role of GenAI in industry transformation, focusing on the following objectives:

1. Understanding GenAI's Technological Foundations: Examining the evolution of generative models and their impact on industry.

2. Assessing Sectoral Applications: Investigating how GenAI is shaping business operations in finance, healthcare, manufacturing, UX design, and project management.

3. Evaluating Economic and Strategic Implications: Analyzing how GenAI enhances productivity, alters business models, and influences labor markets.
4. Addressing Ethical, Regulatory, and Workforce Challenges: Identifying concerns related to AI governance, algorithmic transparency, and job displacement.
5. Developing a Framework for GenAI Adoption: Offering best practices for integrating GenAI into business operations responsibly and sustainably.

Through these objectives, this paper seeks to bridge the gap between technological innovation and real-world industry applications, ensuring that businesses can leverage GenAI effectively while mitigating associated risks.

1.3 Methodology and Theoretical Framework

This research adopts a systematic, interdisciplinary approach that integrates insights from machine learning, engineering management, and business strategy. The methodology includes:

- A comprehensive literature review of existing research on GenAI applications and impacts.
- Case study analysis of businesses and industries that have successfully integrated GenAI.

1.4 Conclusion

Generative AI represents a technological revolution with far-reaching implications across industries. Its ability to augment human capabilities, drive economic growth, and redefine business strategies positions it as a cornerstone of Industry 5.0. However, for businesses to fully capitalize on this transformation, they must navigate the complex landscape of ethical governance, workforce adaptation, and regulatory frameworks.

This paper aims to contribute to the ongoing discourse on AI-driven industry transformation by providing a comprehensive, research-backed framework for understanding, evaluating, and implementing GenAI in business contexts. By fostering responsible AI adoption, industries can harness the full potential of GenAI while ensuring equitable, sustainable, and ethically sound technological advancement.

Chapter 2

GenAI in Business Transformation & GenAI model comparison

2. GenAI in Business Transformation & GenAI model comparison

2.1 How Companies Are Leveraging GenAI for Innovation

Generative AI has emerged as a transformative force in business innovation, unlocking new opportunities across multiple functions. According to McKinsey, approximately 75% of the total annual value from generative AI use cases comes from four areas: marketing, customer service, R&D, and software development (Chui, M., Hazan, E., et al., 2023). These applications are not only enhancing efficiency but also reshaping traditional business models.

Organizations are adopting various strategies to implement GenAI. Some rely on off-the-shelf LLMs, while others customize AI solutions for their specific needs. The choice between pre-trained vs. proprietary models hinges on factors such as data security, regulatory compliance, and industry-specific optimization (Kanbach, D. K., Heiduk, L., et al., 2024). For example, Google and Microsoft are integrating GenAI to automate software development, allowing developers to shift focus from repetitive coding tasks to high-level decision-making (Kanbach, D. K., Heiduk, L., et al., 2024).

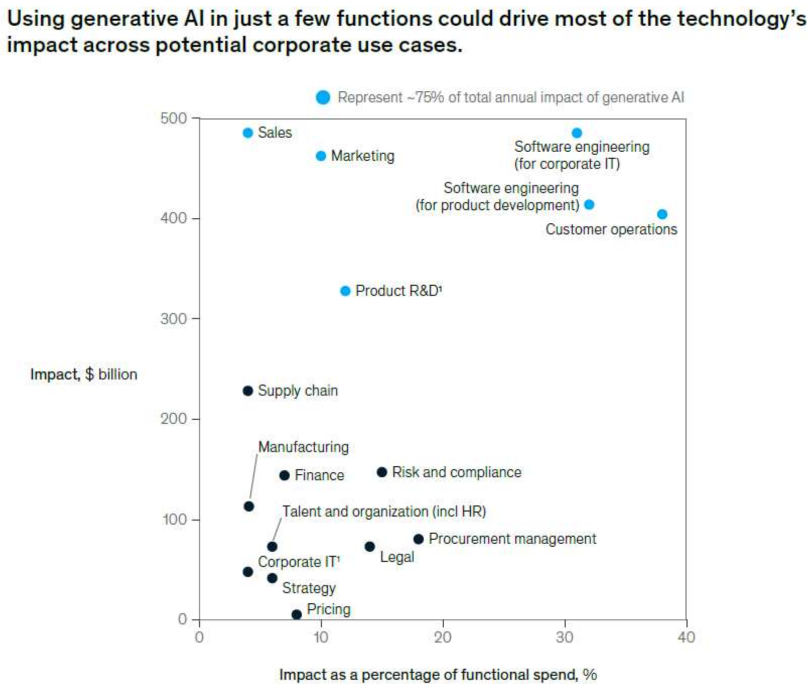

Figure 3 Impact of GenAI[3]

2.2 Applications in Automation, Marketing, Customer Service, and R&D

2.2.1 Marketing and Sales

Marketing has been revolutionized by AI-powered content automation, hyper-personalization, and audience segmentation. AI-driven campaigns can tailor messages to individual preferences, significantly improving customer engagement and conversion rates (GenAI in Business and Industry, 2024). Companies leverage ChatGPT and similar tools to generate advertisements, product recommendations, and email marketing campaigns by analyzing customer behavior patterns (GenAI in Business and Industry, 2024). Moreover, businesses can extract real-time market trends and insights from unstructured data sources such as social media, enhancing customer outreach precision (McKinsey, 2024).

[3] Impact of GenAi Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey & Company, June 2023.

2.2.2 Customer Service and Operations

GenAI is redefining customer service by enabling AI-powered chatbots, digital self-service, and automated support agents. Research indicates that AI-enhanced customer operations can increase issue resolution rates by 14% while reducing handling time by 9% (McKinsey, 2024). Companies such as Lemonade, a leading insurance provider, use AI-driven chatbots to manage customer inquiries in real-time, improving response speed and satisfaction (GenAI in Business and Industry, 2024).

2.2.3 Research & Development (R&D)

Generative AI is accelerating scientific discovery and product innovation. In pharmaceuticals, AI models are revolutionizing drug discovery by analyzing molecular datasets, significantly reducing R&D timelines (McKinsey, 2024). Furthermore, synthetic datasets generated by AI are being used in material sciences and finance, reducing the dependency on real-world experiments and expediting innovation cycles (McKinsey, 2024).

2.3 Examples from Industries Like Finance, Healthcare, and Manufacturing

2.3.1 Finance Industry

In finance, GenAI is transforming fraud detection, risk assessment, and algorithmic trading. AI-powered models analyze vast transactional datasets to detect anomalies, preventing fraudulent activities in real-time (GenAI in Business and Industry, 2024). Additionally, AI-driven financial assistants provide personalized investment recommendations and portfolio optimization, improving customer engagement in the banking sector (GenAI in Business and Industry, 2024).

2.3.2 Healthcare Industry

The healthcare sector is undergoing a profound transformation with AI-driven diagnostics, operational automation, and administrative efficiency. AI models now analyze radiology scans, patient histories, and lab reports, enabling faster and more accurate diagnoses (GenAI in Business and Industry, 2024). Hospitals have also deployed AI-based automation for appointment scheduling and medical record management, reducing administrative burdens on healthcare professionals (GenAI in Business and Industry, 2024).

2.3.3 Manufacturing Industry

Manufacturing benefits from predictive maintenance, supply chain optimization, and production efficiency powered by GenAI. AI systems analyze sensor data from industrial equipment to predict failures before they occur, reducing downtime and improving productivity (GenAI in Business and Industry, 2024). Furthermore, AI-driven logistics optimization ensures efficient inventory management and demand forecasting, minimizing operational costs (GenAI in Business and Industry, 2024).

The adoption of Generative AI across industries signifies a paradigm shift in automation, business strategy, and customer engagement. By leveraging AI-driven models, companies in finance, healthcare, and manufacturing are improving efficiency, reducing costs, and accelerating innovation cycles. However, successful adoption requires addressing key challenges such as data security, model customization, and regulatory compliance. As GenAI continues to evolve, businesses must strategically integrate these technologies to maintain competitiveness in a rapidly transforming market landscape.

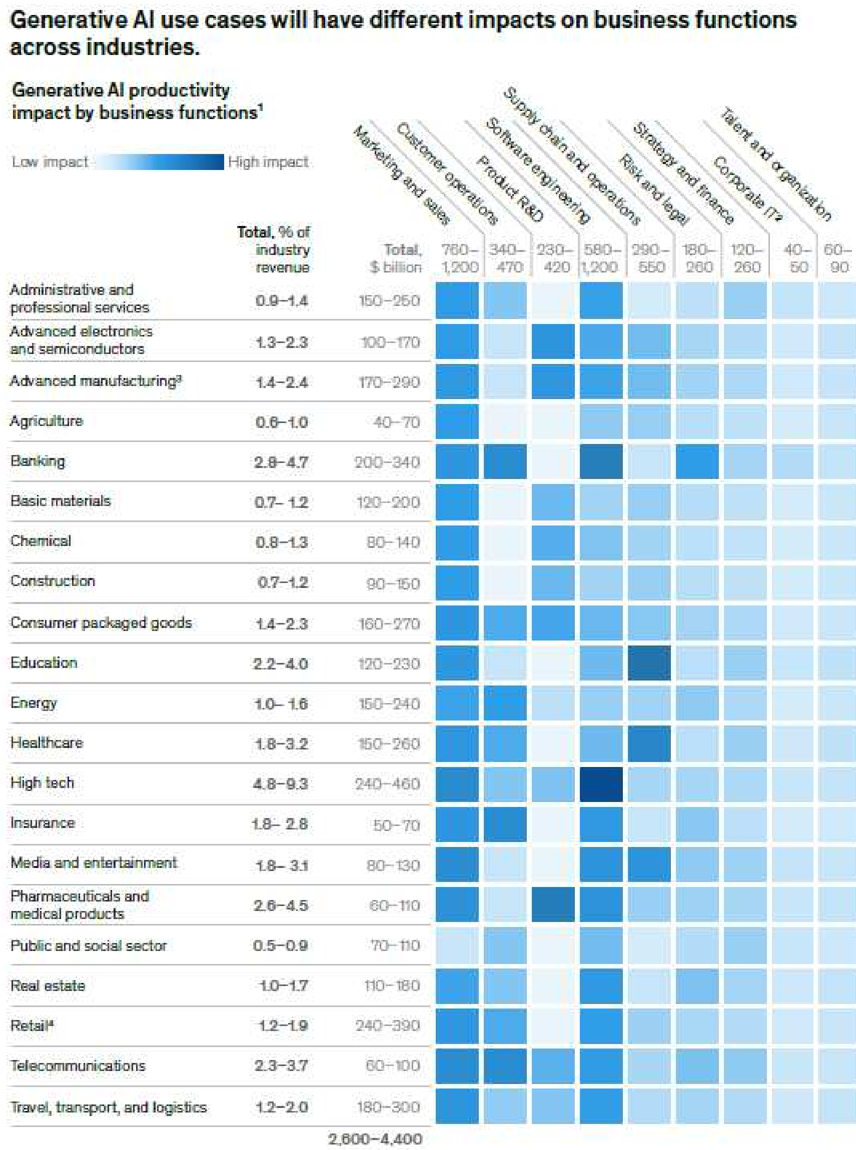

Figure 4 GenAI impact on financial functions[4]

2.4 GenAI opens source and closed source comparison:

[4] GenAI impact on financial functions Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey & Company, June 2023.

The adoption of LLMs in R&D necessitates a systematic approach to ensure optimal utilization of these models. Firstly, selecting an appropriate open-source LLM that aligns with the requirements and resources of the company is crucial. Factors such as model architecture, pre-training data, and computational requirements need to be considered during this selection process. Private models made available through API requests allows for an easier management but with less control over the models. The most important are:

- Claude 3, powerful state of the art model released by Anthropic (04/03/2024) in 3 versions: Haiku, Sonnet and Opus (from the least to the most powerful) .
- ChatGPT, market-leading family of models developed by OpenAI and currently available mainly in 3 versions: GPT-4, GPT-4 Turbo and GPT-3.5 Turbo.
- Gemini, developed by Google and currently available as Gemini 1 Pro, with the more powerful Gemini 1 Ultra and the recent Gemini 1.5 Pro (15/02/2024) both available in preview
- Command, made by Cohere and available in 3 versions: Command Light, Command and the recent addition Command-R optimized for RAG (11/03/2024, also available freely for research purposes).
- Titan, developed by Amazon as a core model for Bedrock and available in 2 versions: Light and Express.
- Jurassic 2, developed by AI21 labs id offered in 3 versions: Light, mid and Ultra.
- Mistral, developed by MistralAI and available in 3 versions: Small, Medium and Large.
- On the contrary, Open-source models, result to be much more flexible but less powerful:
- Llama 2, family of models developed by Meta and available in 3 sizes: 7B, 13B and 70B.
- Mistral 7B and Mixtral 8x7B, the two open models developed by MistralAI.
- Gemma, family of light models recently released by Google (21/02/2024) and available in 2 sizes: 2B and 7B.

2.4.1 Model Parameters

Model parameters are the trainable weights of a neural network. In the context of LLMs, they determine how the model "learns" language, patterns, and relationships from data. Key aspects include:

- Parameter Count:

  The total number of weights (often in the billions) that a model uses. A higher count often suggests increased capacity for learning complex patterns, though it also demands more computational resources.

- Architecture Details:

  This includes the number of layers (depth), hidden layer sizes (width), attention heads, and other hyperparameters that define the model's structure. These settings influence both performance and efficiency.

- Vocabulary and Tokenization:

  The set of tokens (words, sub words, or characters) the model recognizes. A larger vocabulary can help the model handle diverse linguistic inputs.

- Training Data:

  The diversity and volume of text used during training, which directly affect the model's generalization and performance on various tasks.

Deciding to focus on open source model, Table 1 reports some important characteristics (namely number of parameters per model and relative VRAM necessary for usage, pre-training size, context window size, type of license and commercial use availability) of the most common LLMs.

The following characteristics are evaluated and compared.

1. Number of Parameters: the total number of weights or learnable values in a model. These parameters define the model's capacity to understand and generate language. More parameters generally increase a model's ability to perform complex tasks, understand nuances, and generate high-quality outputs. However, larger models require more computational resources.

2. GPU/RAM Required (fp/bf): the amount of GPU memory (VRAM) and RAM needed to run a model, typically expressed in terms of floating-point precision (fp) or brain-float (bf) precision. "fp16" (half-precision) and "bf16" are common ways to reduce memory requirements while maintaining performance. Higher model sizes demand more memory for processing. Models using lower precision (fp16/bf16) can run faster and use less GPU memory, making them more efficient for deployment.

3. Pre-train Size (Tokens): The number of tokens (chunks of words or characters) used in the initial training phase, before fine-tuning. This is how much data the model is exposed to while learning general language patterns. Larger pre-training sizes allow the model to learn from more diverse and extensive data, improving its performance on a wider range of tasks. A bigger pre-train size often leads to a more generalizable model.

4. Context Window: the maximum number of tokens the model can process and consider at once when generating a response. It represents the model's ability to "remember" the input during processing. A larger context window allows the model to handle longer texts and maintain coherence in conversations or documents. It's crucial for tasks that require understanding large amounts of context or for longer interactions.

5. License: the legal agreement that defines how a model can be used, shared, and modified. This could range from open-source licenses (e.g., MIT, Apache) to commercial licenses or proprietary licenses. The license impacts the freedom and limitations of how a model can be deployed, shared, or adapted for specific use cases, especially in business or research. It helps protect intellectual property and ensures legal usage.

6. Free Commercial Use: Whether the model's license permits its use for commercial purposes without needing to pay royalties or ask for permission. Free commercial use makes the model accessible for

businesses and startups to incorporate into their products or services. If a model requires a commercial license or payment, it limits who can use it commercially.

MODEL	Number of Parameters	GPU/RAM required (fp/bf)	Pre-train size (tokens)	Context Window	License	Free Commercial Use
Gemma 2B	2.51B	8GB	2 trillions	8k	Gemma License	YES
Gemma 7B	8.54B	24GB	6 trillions	8k	Gemma License	YES
Llama 2 7B	7B	16GB	2 trillions	4k	Llama 2 Community License	YES (limit of 700 mil. users)
Llama 2 13B	13B	32GB	2 trillions	4k	Llama 2 Community License	YES (limit of 700 mil. users)
Llama 2 70B	70B	145GB	2 trillions	4k	Llama 2 Community License	YES (limit of 700 mil. users)
Mistral 7B	7.3B	16GB	N/A	32k	Apache 2.0	YES
Mistral 8x7B	46.7B (12.9B active)	100GB	N/A	32k	Apache 2.0	YES
DeepSeek-R1	671B (37B Active)		14.8 trillion	131K	MIT License	YES

Table 1 Open source LLM characteristics

With the aim to select one open source LLM to implement the tuning and the product analysis, many indices developed to evaluate language models can be adopted. Table 2 report the models' performances, evaluated with the aid of various popular benchmark commonly used, focusing in particular on:

- Knowledge capabilities and knowledge retention, using benchmarks like MMLU, ARC-e and GPQA.
- Text comprehension and reasoning over text, using benchmarks like DROP, HotPotQA, TriviaQA and Kilt.

- Common sense and reasoning, using benchmarks like HellaSwag, Big-Bench-Hard and WinoGrande.
- General perception and usability, using benchmark like Chatbot Arena (ranked elo system obtained through anonymous “duels” evaluated by humans).

Note that not all measures are available for all models and that they have been chosen to evaluate also commercial models (not reported here) and the most performing model on each relevant measure are highlighted.

The definitions and details of each adopted benchmark measure are reported hereafter.

1. **GPQA: A Graduate-Level Google-Proof Q&A Benchmark** : A benchmark designed to test high-level reasoning abilities in AI models. The questions are challenging, with PhD-level experts achieving 65% accuracy (74% after corrections), while skilled non-experts score only 34% despite unrestricted web access. Even advanced models like GPT-4 struggle, reaching just 39% accuracy. The dataset is intentionally "Google-proof," highlighting the necessity for scalable oversight methods to ensure the reliability of AI-generated scientific insights, especially when human supervisors struggle to evaluate them.

   Figure 5 reports the level of such index for different up to date models.

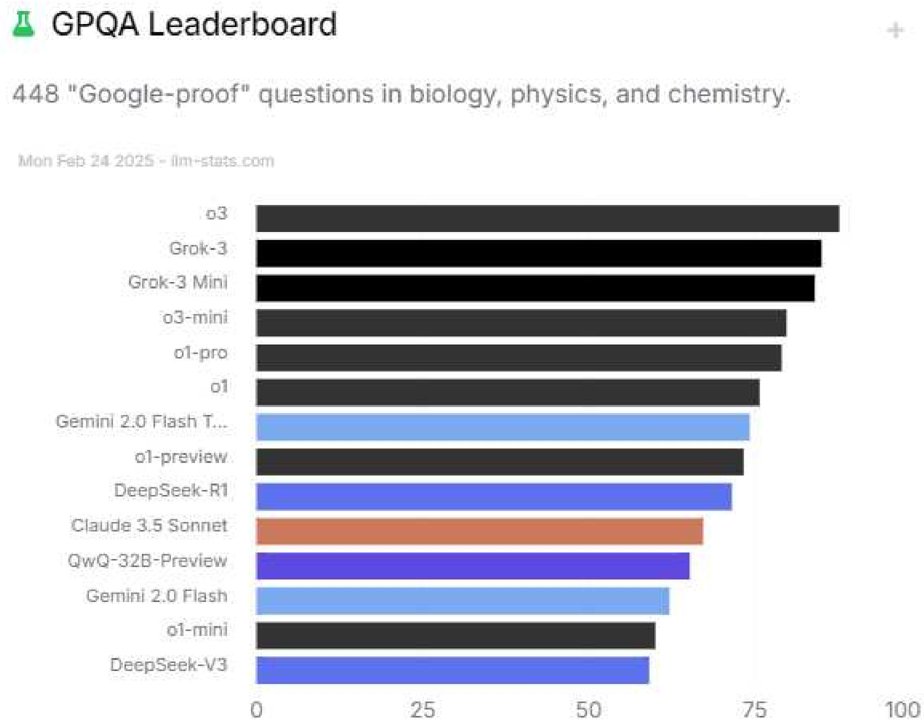

Figure 5 GPQA leaderboard comparison[5]

2. **Measuring Massive Multitask Language Understanding:**

   This benchmark evaluates the multitask accuracy of language models across 57 diverse subjects, including mathematics, history, law, and computer science. It requires extensive world knowledge and problem-solving capabilities. While smaller models perform at near-random levels, the largest GPT-3 model surpasses this by 20 percentage points. However, no model has yet achieved expert-level accuracy, and performance remains inconsistent, particularly in critical fields such as morality and law. This benchmark is instrumental in assessing the academic and professional understanding of AI systems while identifying their key limitations.

[5] GPQA leaderboard comparison From internet "https://llm-stats.com/"

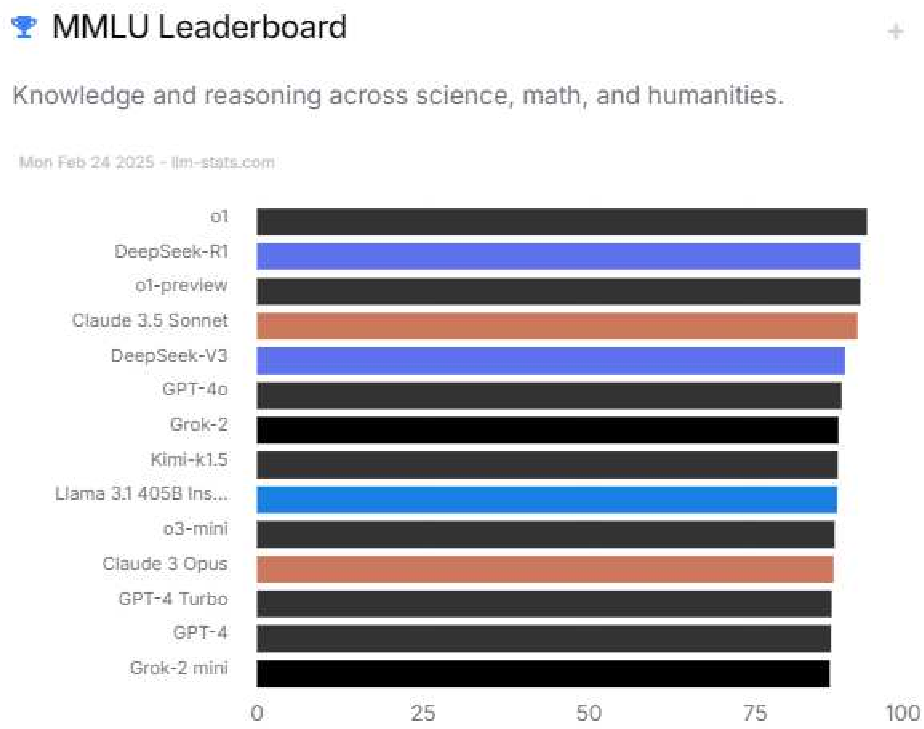

Figure 6 MMLU leader board comparison[6]

3. The Concept ARC Benchmark: Evaluating Understanding and Generalization in the ARC Domain:

Abstraction and Reasoning Corpus (ARC) and is designed to measure AI's ability to abstract and generalize across spatial and semantic concepts. Unlike the original ARC dataset, Concept ARC organizes problems based on specific concepts with varying complexity levels. Testing on humans, top ARC competition solvers, and GPT-4 demonstrates that humans vastly outperform AI in conceptual reasoning, emphasizing the existing gap in AI's abstraction capabilities. This benchmark aims to advance AI research in conceptual understanding and evaluation.

[6] MMLU leader board comparison From internet "https://llm-stats.com/"

4. **HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering**

   A dataset comprising 113,000 question-answer pairs sourced from Wikipedia, designed to test multi-hop reasoning capabilities. The dataset features:

   - Questions that require synthesizing information from multiple supporting documents.
   - Diversity in question formulation, not constrained to pre-existing knowledge bases.
   - Sentence-level supporting facts to enable reasoning under strong supervision.
   - Factoid comparison questions that assess an AI system's ability to extract relevant facts and perform logical comparisons.

   Evaluations reveal that HotpotQA remains a challenging dataset for modern QA systems, though the provided supporting facts significantly improve explainability and performance.

5. **KILT: a Benchmark for Knowledge Intensive Language Tasks**

   A benchmark for tasks such as open-domain question answering, fact-checking, entity linking, and slot filling. It standardizes task evaluation by grounding all tasks in the same Wikipedia snapshot, reducing the need for customized knowledge indexing and enabling the development of more generalizable AI models. Findings indicate that a shared dense vector index, coupled with a sequence-to-sequence model, outperforms more specialized approaches in fact-checking, open-domain QA, and dialogue generation.

6. **HellaSwag: Can a Machine Really Finish Your Sentence?**

   A benchmark designed to evaluate commonsense natural language inference by presenting models with ambiguous event descriptions. Unlike prior commonsense inference tasks where models achieved near-human performance, HellaSwag remains a challenge for state-of-the-art systems, achieving less than 48% accuracy compared to human performance above 95%. The dataset is constructed using Adversarial Filtering (AF), ensuring that generated wrong answers are challenging for AI while remaining obvious to humans. This highlights the limitations of deep learning models in commonsense reasoning and suggests an adversarial approach to improving NLP benchmarks.

7. **BIG-Bench Extra Hard:**

An extension of the BIG-Bench dataset, designed to evaluate the general reasoning abilities of large language models (LLMs). As recent advancements have saturated BIG-Bench and BIG-Bench Hard (BBH), BBEH introduces new, more difficult tasks probing similar reasoning skills. Evaluations show that the best general-purpose models achieve an accuracy of only 9.8%, while reasoning-specialized models score 44.8%, indicating significant room for improvement in AI's general reasoning capabilities.

8. **WinoGrande :**

An extension of the Winograd Schema Challenge (WSC), consisting of 44,000 commonsense reasoning problems. It was developed to address biases in WSC while increasing scale and difficulty. The dataset construction process included:

- A structured crowdsourcing approach for diverse problem generation.
- Bias reduction using the AfLite algorithm to eliminate word associations exploitable by AI models.

State-of-the-art methods achieve 59.4-79.1% accuracy, significantly below human performance at 94%. Results suggest that while WinoGrande facilitates transfer learning across commonsense benchmarks, it also underscores the potential overestimation of machine commonsense abilities due to dataset biases.

9. **Chatbot Arena**:

A human-evaluated ranking system using an Elo score obtained through anonymous "duels" between chatbots. This benchmark assesses general perception and usability, helping measure LLM effectiveness in open-ended conversational AI tasks.

MODEL	MMLU	ARC-e	GPQA	Drop	HotPotQA	KILT	TriviaQA	HellaSwag	WinoGrand e	Big-Bench-Hard	ChatBOT Arena
Gemma 2B	42.3%	73.2%	-	-	-	-	53.2%	71.4%	65.4%	35.2%	985
Gemma 7B	64.3%	81.5%	-	-	-	-	63.2%	81.2%	72.3%	55.1%	1029
Llama 2 7B	44.4%	68.7%	-	-	-	-	56.6%	77.2%	69.5%	32.6%	1027
Llama 2 13B	55.6%	75.2%	-	-	-	-	64,0%	80.7%	72.9%	39.4%	1043
Llama 2 70B	69.9%	79.9%	-	-	42%	63.2%	73.0%	85.4%	80.4%	51.2%	1082
Mistral 7B	62.5%	80.5%	-	-	-	-	62.5%	81.0%	74.2%	56.1%	1073
Mixtral 8x7B	70.6%	83.6%	-	-	56%	66.0%	71.5%	84.4%	77.2%	-	1114

Table 2 Open source LLM comparison

2.4.1.1 LLMs: commercial and open-source products evaluation

In recent R&D environments, the systematic evaluation of LLMs is essential to ensure their optimal application. An initial step involves choosing an open-source LLM that aligns with the specific requirements and resources of an organization. This selection process should consider factors such as model architecture, pre-training data, and computational demands. In contrast, proprietary models offered via API provide streamlined management but at the cost of reduced control over the model's inner workings. Key commercial models include:

- DeepSeek-R1 is the first-generation reasoning model built atop DeepSeek-V3 (671B total parameters, 37B activated per token). It incorporates large-scale reinforcement learning (RL) to enhance its chain-of-thought and reasoning capabilities, delivering strong performance in math, code, and multi-step reasoning tasks.
- Claude 3 by Anthropic, a state-of-the-art system released on 04/03/2024 in three variants—Haiku, Sonnet, and Opus (from least to most powerful).
- ChatGPT from OpenAI, currently available in three primary versions: GPT-4, GPT-4 Turbo, and GPT-3.5 Turbo.
- Gemini from Google, available as Gemini 1 Pro, with the more powerful Gemini 1 Ultra and the preview version Gemini 1.5 Pro released on 15/02/2024.
- Command by Cohere, offered in three versions—Command Light, Command, and the recently introduced Command-R (optimized for RAG and also available freely for research purposes as of 11/03/2024).
- Titan from Amazon, a core model for Bedrock available in two versions: Light and Express.
- Jurassic 2 by AI21 Labs, provided in three variants: Light, Mid, and Ultra.
- Mistral from Mistral AI, available in three sizes: Small, Medium, and Large.

On the other hand, open-source models tend to offer greater flexibility, though they generally deliver lower performance relative to their proprietary counterparts. Notable examples include:

- Llama 2 by Meta, available in three sizes: 7B, 13B, and 70B.

- Mistral 7B and Mixtral 8x7B, both developed by Mistral AI.
- Gemma, a family of lightweight models recently released by Google on 21/02/2024, available in 2B and 7B sizes.

LLM comparisons across benchmark prices

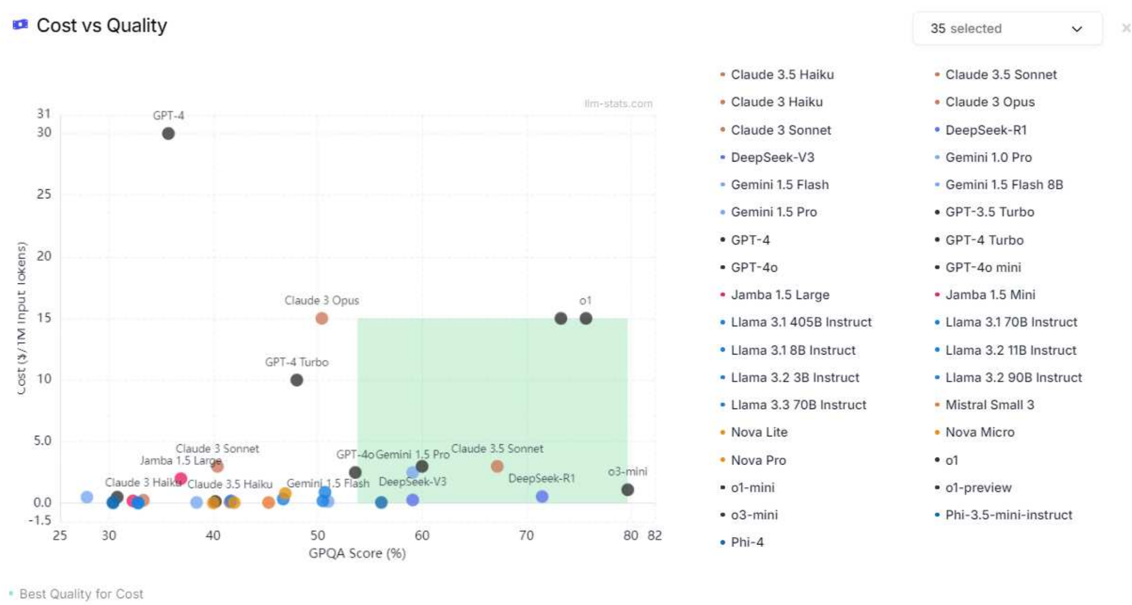

Figure 7 Cost vs Quality comparison[7]

[7] Cost vs Quality comparison From internet “https://llm-stats.com/”

Chapter 3

Opportunities and Economic Potential of GenAI

3 Opportunities and Economic Potential of GenAI

3.1 Productivity and Economic Growth Potential

The transformative potential of Generative AI extends far beyond its immediate technological applications, offering a significant economic catalyst for global productivity growth. According to McKinsey, the automation of individual work activities enabled by GenAI could provide the global economy with an annual productivity boost of 0.5% to 3.4% from 2023 to 2040, with GenAI alone contributing 0.1% to 0.6% annually (Chui, M., Hazan, E., et al., 2023). Such gains, while modest in absolute percentage points, have compounding long-term effects on economic growth, especially in economies facing demographic challenges such as aging populations and slowing workforce expansion.

One of the driving forces behind these productivity gains is the ability of GenAI to reduce the time spent on repetitive tasks, enabling human capital to be allocated to higher-value functions. This shift can be seen in industries where labor shortages and skill gaps threaten economic sustainability. McKinsey further emphasizes that labor productivity will need to accelerate to maintain economic prosperity in countries experiencing shrinking workforces due to demographic shifts (Chui, M., Hazan, E., et al., 2023).

However, for these productivity benefits to materialize, substantial investments must be made in worker upskilling and labor redeployment. McKinsey highlights that Generative AI will require investments in workforce transition strategies, ensuring that displaced workers are repositioned into roles that complement AI rather than compete with it (Chui, M., Hazan, E., et al., 2023). Thus, while GenAI has the potential to reshape labor markets, proactive management will be essential to maximizing its benefits while mitigating social and economic disruptions.

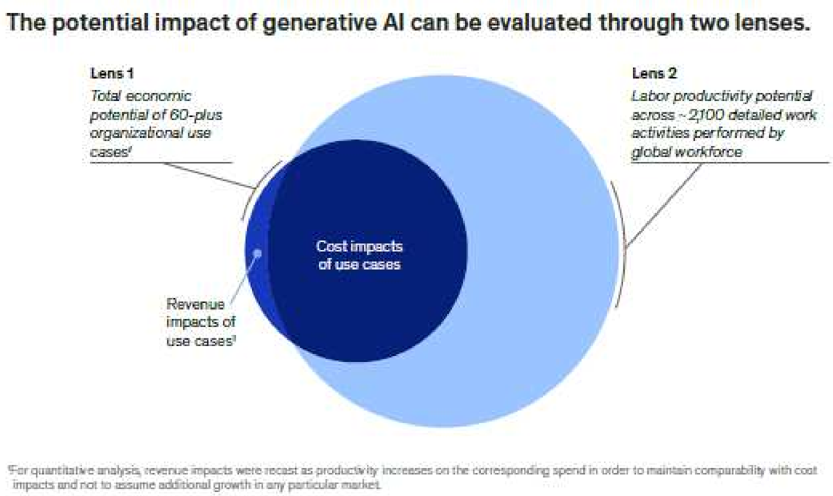

Figure 8 The economical potential impact of GenAI[8]

3.1.1 Impact on Revenue Models and Cost Structures

Beyond productivity enhancements, Generative AI is poised to revolutionize business revenue models and cost structures. As AI systems become more sophisticated, the cost of generating high-quality content, automated processes, and real-time insights is drastically decreasing, allowing businesses to explore novel monetization strategies and pricing structures (Kanbach, D. K., Heiduk, L., et al., 2024).

A critical example is the freemium model, which has gained significant traction with the rise of AI-generated content. Companies can now offer basic AI-driven services for free while monetizing premium features such as advanced content personalization, enhanced automation, or enterprise-grade capabilities (Kanbach, D. K., Heiduk, L., et al., 2024). This shift has created new opportunities for businesses to increase customer engagement and long-term revenue streams without incurring substantial marginal costs.

Moreover, Generative AI enables greater cost efficiency in content-driven industries, such as digital media, advertising, and entertainment. According to McKinsey, the total economic benefits of Generative AI—including major use cases and productivity enhancements—amount to $6.1 trillion to $7.9 trillion annually,

[8] The economical potential impact of GenAI GenAI impact on financial functions Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey & Company, June 2023.

a staggering figure that underscores the vast financial implications of widespread AI adoption (Chui, M., Hazan, E., et al., 2023).

Companies that strategically deploy GenAI are finding that it not only reduces operational costs but also creates new value streams, unlocking entirely new business paradigms based on hyper-personalization, automation, and scalable intelligence (Kanbach, D. K., Heiduk, L., et al., 2024). For example, AI-driven product recommendations, real-time customer insights, and autonomous customer service interactions are reshaping traditional revenue models, driving efficiency and profitability across industries.

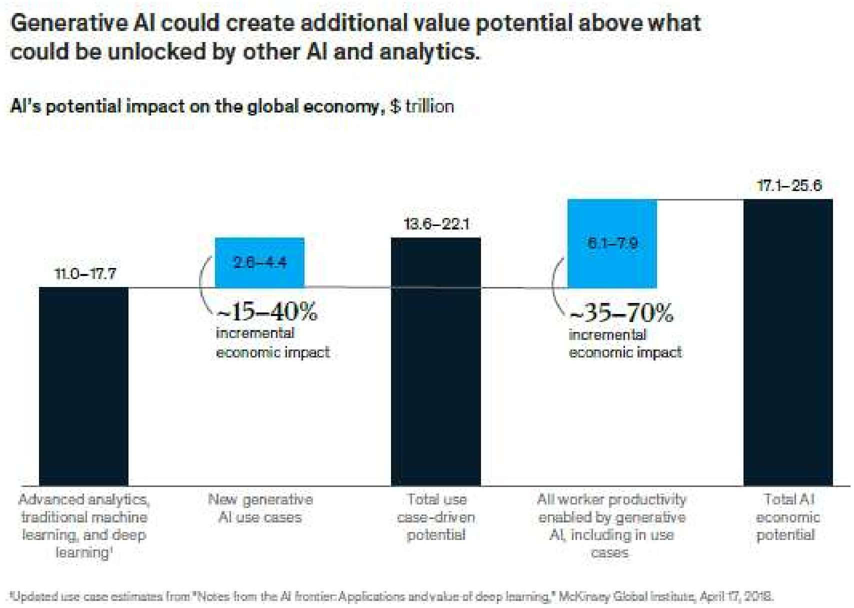

Figure 9 AI's potential impact on the global economy[9]

3.1.2 Sector-Wise Breakdown of GenAI's Contribution

The economic impact of Generative AI is not evenly distributed across industries; rather, certain sectors are experiencing outsized benefits from AI-driven efficiencies and new business opportunities. Among the

[9] AI's potential impact on the global economy GenAI impact on financial functions Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey & Company, June 2023.

most affected industries are finance, retail, high-tech, and life sciences, where AI-driven automation, decision-making, and predictive analytics are generating substantial economic value.

The financial sector stands to gain significantly from Generative AI, with potential value creation estimated at $200 billion to $340 billion annually (Chui, M., Hazan, E., et al., 2023). Key areas where Generative AI is driving value include fraud detection, algorithmic trading, automated compliance monitoring, and AI-powered financial advising. These AI-driven advancements are not only reducing operational costs but also improving risk assessment and decision-making capabilities, ensuring greater regulatory compliance and financial security.

Retail and consumer goods industries are projected to experience an economic impact of $400 billion to $660 billion annually due to AI-driven efficiencies (Chui, M., Hazan, E., et al., 2023). The primary drivers of this growth include AI-powered demand forecasting, personalized shopping experiences, automated supply chain management, and real-time inventory optimization. As consumer behavior becomes increasingly data-driven, retailers are leveraging AI to offer dynamic pricing, AI-generated marketing content, and hyper-personalized shopping recommendations.

Industries at the cutting edge of innovation—such as high-tech and life sciences—are also among the largest beneficiaries of Generative AI. The ability to analyze vast datasets at scale is accelerating drug discovery, biotech research, and personalized medicine, leading to breakthrough treatments and new business opportunities (Chui, M., Hazan, E., et al., 2023). Similarly, the high-tech sector is utilizing AI to enhance software development, optimize cloud computing operations, and automate complex decision-making processes.

The economic opportunities presented by Generative AI are immense and transformative, spanning productivity growth, revenue model innovation, and sector-specific advancements. With an estimated annual impact ranging from $6.1 trillion to $7.9 trillion, GenAI is emerging as a defining force in the global economy (Chui, M., Hazan, E., et al., 2023). However, realizing these benefits requires strategic investments in AI adoption, workforce training, and policy frameworks to navigate the disruptions that AI-driven automation may introduce.

Industries such as finance, retail, high-tech, and life sciences are already capitalizing on these advancements, leveraging AI-driven efficiencies to drive revenue growth, enhance customer experiences, and optimize decision-making. As AI adoption accelerates, organizations that proactively integrate AI into their operational and strategic frameworks will be best positioned to thrive in this new era of economic transformation.

3.2 GenAI in Operations & Risk Management

The shift from Industry 4.0 to Industry 5.0 marks a significant transition from automation-focused manufacturing to a more human-centric and sustainable industrial framework. While Industry 4.0 emphasized digital transformation, interconnected devices, and automation, Industry 5.0 seeks to integrate human creativity, resilience, and sustainability into industrial processes (The Role of Generative AI in Managing Industry Projects, 2024).

One of the defining characteristics of Industry 5.0 is collaboration between humans and AI systems, ensuring that automation does not replace human input but rather enhances decision-making and adaptability. Industry leaders argue that the ability of AI to optimize workflows, reduce errors, and provide predictive insights will allow workers to focus on high-value, creative problem-solving tasks (The Role of Generative AI in Managing Industry Projects, 2024). This paradigm shift necessitates the integration of AI into manufacturing, logistics, and operations, where real-time data analytics can optimize industrial efficiency (The Role of Generative AI in Managing Industry Projects, 2024).

Furthermore, resilience and sustainability are core principles of Industry 5.0, as highlighted by the European Commission. Generative AI is being integrated to improve resource utilization, enhance waste management, and drive energy efficiency within industrial production (The Role of Generative AI in

Managing Industry Projects, 2024). This transition ensures that industrial ecosystems become more adaptable and sustainable, paving the way for a more responsible AI-driven future.

The integration of Generative AI into operations and risk management marks a pivotal evolution in industrial efficiency and strategic decision-making. As industries transition from Industry 4.0 to Industry 5.0, the emphasis is shifting from pure automation to human-AI collaboration, fostering resilience, creativity, and sustainability.

3.2.1 Supply Chain Optimization and Automation

The role of Generative AI in supply chain optimization is becoming increasingly prominent, with AI-powered systems streamlining inventory management, demand forecasting, and logistics planning (GenAI in Business and Industry, 2024). One of the key applications of AI in supply chains is predicting demand fluctuations with high accuracy, allowing businesses to optimize inventory levels and reduce holding costs (GenAI in Business and Industry, 2024).

Moreover, real-time data analytics has transformed supply chain resilience, enabling companies to quickly adapt to disruptions caused by geopolitical events, market fluctuations, or unexpected demand surges (GenAI in Business and Industry, 2024). By leveraging AI-powered risk assessment tools, businesses can create agile supply chains that minimize delays and ensure operational continuity.

Automation also extends to warehouse management and logistics, where AI-driven robotics and automation systems are streamlining operations. AI-powered autonomous robots are increasingly used for inventory management, picking, and packing, leading to enhanced efficiency and fewer human errors in warehouse operations (GenAI in Business and Industry, 2024). While these advancements have significantly improved accuracy and productivity, challenges such as data privacy concerns, cybersecurity threats, and ethical implications of AI-driven decision-making in logistics remain key considerations (GenAI in Business and Industry, 2024).

3.2.2 Risk Management and Decision-Making Support

Generative AI plays a critical role in risk management and strategic decision-making by enabling companies to analyze large datasets, forecast risks, and implement preventive measures. AI-powered decision-support systems assist in identifying potential threats before they escalate into crises, allowing businesses to develop contingency plans and risk mitigation strategies (The Role of Generative AI in Managing Industry Projects, 2024).

For instance, AI-driven models are being used to support contingency planning by assessing financial risks, operational vulnerabilities, and market uncertainties (The Role of Generative AI in Managing Industry Projects, 2024). AI can generate detailed risk reports, propose response strategies, and analyze the impact of different risk scenarios, helping organizations navigate volatile environments more effectively.

Moreover, AI is revolutionizing procurement and resource management by developing intelligent procurement plans, supplier evaluation frameworks, and contract negotiation strategies (The Role of Generative AI in Managing Industry Projects, 2024). The ability of AI to analyze supply chain risks, detect potential bottlenecks, and recommend alternative sourcing strategies makes it a powerful tool in project and operations management.

Beyond risk management, AI-driven decision-support systems are increasingly used for long-term strategic planning. Companies are leveraging Generative AI to improve forecasting accuracy, optimize resource allocation, and enhance crisis response (The Role of Generative AI in Managing Industry Projects, 2024). By integrating AI into decision-making processes, industries can significantly improve agility and resilience, ensuring that their operations remain robust even in highly dynamic and unpredictable environments.

3.3 Challenges of Implementing GenAI in Industry

3.3.1 Organizational Barriers: Skill Gaps and Infrastructure Readiness

The widespread adoption of Generative AI (GenAI) presents substantial organizational challenges, particularly in terms of workforce skill gaps and infrastructure readiness. The integration of AI systems demands that employees develop new capabilities, shifting from traditional workflows to AI-augmented processes. However, many employees lack the technical proficiency to leverage AI-driven tools effectively, creating a pressing need for corporate upskilling initiatives (GenAI in Business and Industry, 2024).

According to recent studies, organizations must invest in comprehensive upskilling and reskilling programs to address the growing demand for technical competencies such as data science, AI model training, and prompt engineering (GenAI in Business and Industry, 2024). Additionally, as AI automates routine tasks, employees must cultivate soft skills such as adaptability, problem-solving, and collaboration to remain relevant in a GenAI-enabled workforce.

Beyond employee readiness, leadership support is crucial for AI adoption. Research suggests that without strong executive backing, AI initiatives often stagnate, leading to fragmented adoption across departments (The Transformative Potential of Generative AI, 2024). For AI implementation to succeed, organizations must foster a culture that embraces innovation, encouraging employees to experiment with AI-driven solutions while providing necessary training and resources.

From an infrastructure standpoint, IT scalability and system interoperability remain significant challenges. Many companies operate on legacy systems that are not designed to integrate seamlessly with AI-driven workflows, requiring costly upgrades or full-scale replacements (The Transformative Potential of Generative AI, 2024). Ensuring scalability to accommodate increasing AI workloads and expanding data volumes is essential for businesses looking to implement AI solutions effectively.

As enterprises progress in their digital transformation journeys, they must navigate these organizational and technical challenges by investing in training, leadership-driven AI adoption strategies, and modernized IT infrastructures.

3.3.2 Security Risks and Data Privacy Concerns

Security and privacy remain at the forefront of GenAI adoption concerns. The handling of sensitive customer and corporate data by AI-driven systems raises significant compliance and security risks. In particular, industries such as finance and healthcare, which manage large volumes of personally identifiable information (PII), face heightened challenges in ensuring AI-driven processes comply with privacy regulations (Generative AI in UX Design and Research, 2024).

One of the most pressing concerns is GenAI's vulnerability to cyberattacks. Research highlights that AI-powered tools are susceptible to malicious exploitation, including AI-generated phishing campaigns, automated malware creation, and deepfake-driven fraud (The Transformative Potential of Generative AI, 2024). Furthermore, cybercriminals now leverage AI-enhanced hacking tools to breach enterprise security systems, making traditional cybersecurity measures inadequate against evolving AI threats.

Beyond external threats, insider risks associated with AI usage must also be considered. Without proper governance structures in place, employees may inadvertently expose sensitive data to AI models that lack adequate security protocols (Generative AI in UX Design and Research, 2024). Thus, organizations must establish AI governance frameworks to oversee how AI systems process and store data, ensuring compliance with regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

To mitigate these risks, businesses must prioritize cybersecurity investments and AI governance strategies. Key measures include limiting AI access to sensitive datasets, employing advanced encryption protocols, and implementing continuous AI model auditing. By proactively addressing these security vulnerabilities, enterprises can maximize the benefits of GenAI while safeguarding critical data assets.

3.3.3 Lack of Formal Policies and Regulations

Despite the rapid proliferation of Generative AI, most companies lack formalized policies governing its use. Studies indicate that many employees use GenAI tools without clear directives from management, resulting in inconsistent adoption and potential compliance violations (Generative AI in UX Design and Research, 2024).

The regulatory landscape surrounding GenAI is still evolving, with governments and industry regulators struggling to keep pace with AI advancements. This regulatory lag creates significant challenges, particularly in sectors where AI-driven decisions carry legal and ethical implications. For example, GenAI applications in content creation, automated decision-making, and algorithmic hiring pose risks related to copyright infringement, discrimination, and misinformation (The Transformative Potential of Generative AI, 2024).

Intellectual property (IP) concerns further complicate the regulatory environment. AI-generated content raises legal questions about copyright ownership, as many AI models are trained on publicly available data without explicit licensing agreements (The Transformative Potential of Generative AI, 2024). As legal frameworks remain ambiguous, businesses risk legal disputes over the unauthorized use of copyrighted materials in AI-generated outputs.

Regulatory agencies have classified GenAI as a high-risk technology, necessitating stringent compliance measures. However, the lack of clear guidelines has resulted in prolonged approval processes for AI-driven products, delaying time-to-market for innovations (Generative AI in UX Design and Research, 2024). Policymakers must accelerate the development of AI-specific regulations to address these concerns while balancing innovation with ethical and legal safeguards.

In response to these regulatory gaps, organizations must proactively develop internal AI policies that define acceptable AI applications, ethical AI usage principles, and compliance protocols. Establishing internal AI governance committees can ensure responsible AI deployment while mitigating potential legal and reputational risks.

The implementation of Generative AI in industry is fraught with organizational, security, and regulatory challenges. Addressing workforce skill gaps, IT infrastructure limitations, and leadership support deficiencies is paramount for companies looking to harness AI's full potential. Without a strategic approach to AI upskilling and system integration, businesses may struggle to realize GenAI's transformative benefits.

Security and privacy risks pose additional obstacles, necessitating robust AI governance frameworks, data protection policies, and cybersecurity investments. Companies must remain vigilant against AI-driven cyber threats while ensuring compliance with data protection laws and ethical AI guidelines.

Furthermore, the absence of comprehensive AI regulations has created legal uncertainties around copyright, compliance, and algorithmic decision-making. As policymakers work to establish clearer AI governance structures, organizations must take proactive steps to define their own AI policies, ensuring responsible and legally sound AI implementations.

By addressing these challenges, businesses can navigate the complexities of GenAI adoption while maximizing its potential for innovation, efficiency, and long-term competitiveness.

3.3.4 Legal, Ethical, and Societal Challenges

The rapid integration of Generative AI into business and society presents both unprecedented opportunities and formidable challenges. While GenAI enhances efficiency, automation, and economic productivity, it also raises serious concerns about algorithmic bias, misinformation, labor displacement, and intellectual property conflicts.

As Generative AI continues its rapid expansion, it raises ethical, regulatory, and societal challenges that must be carefully navigated. One of the primary concerns is algorithmic bias, which arises from the reliance of AI models on extensive datasets that may include historical biases, discriminatory patterns, or incomplete representations of diverse populations. Research has highlighted that GenAI models trained on biased data can perpetuate and even amplify these biases, leading to inequitable outcomes in areas such as hiring, lending decisions, and law enforcement (GenAI in Business and Industry, 2024).

Moreover, AI-generated misinformation is a growing concern, particularly in media and communications. The proliferation of deepfakes and AI-generated text, images, and videos has raised serious concerns regarding misinformation, public manipulation, and trust erosion (GenAI in Business and Industry, 2024). These risks are exacerbated by the ability of GenAI to generate convincing yet false narratives, potentially influencing public opinion and electoral processes.

Given these challenges, there is a pressing need for robust AI governance frameworks that promote transparency, fairness, and accountability in the deployment of GenAI systems (GenAI in Business and Industry, 2024). The establishment of regulatory guidelines that mandate AI auditing, explainability, and ethical oversight is essential to mitigate the risks associated with biased decision-making and misinformation.

The adoption of Generative AI is transforming the labor market, raising questions about job displacement, workforce reskilling, and the future of employment. Analysts predict that AI will have the highest impact on white-collar knowledge workers, particularly in roles requiring advanced language and communication skills (Kanbach, D. K., Heiduk, L., et al., 2024).

According to reports, 300 million jobs could be affected by AI, with up to 40% of all working hours impacted by LLMs like GPT-4 (Kanbach, D. K., Heiduk, L., et al., 2024). This transformation is expected to affect occupations such as data entry, customer service, paralegal work, and financial analysis, where AI can perform repetitive and rule-based tasks with higher efficiency and accuracy.

While concerns about job displacement are valid, AI adoption is also creating new employment opportunities. The rise of AI-driven industries has led to an increase in job postings related to AI development, model fine-tuning, and AI ethics research (Kanbach, D. K., Heiduk, L., et al., 2024). Additionally, in care-oriented professions, AI can free up human labor by automating administrative tasks, allowing workers to focus on human-centric roles such as elder care, nursing, and therapy (Kanbach, D. K., Heiduk, L., et al., 2024).

To mitigate job losses, governments and businesses must invest in workforce transition programs, equipping employees with AI literacy, prompt engineering skills, and interdisciplinary AI applications knowledge. The reallocation of labor into emerging AI-driven roles will be crucial in ensuring economic stability and sustainable employment in the AI era.

One of the most contentious debates surrounding Generative AI is its impact on intellectual property (IP) rights and copyright laws. AI systems ingest vast amounts of publicly available data to generate new content, raising questions about whether the outputs produced by GenAI models constitute derivative works of copyrighted material.

A landmark ruling by the U.S. Copyright Office and the U.S. District Court for the District of Columbia affirmed that AI-generated works without human intervention are not eligible for copyright protection (The Role of Generative AI in Managing Industry, 2024). This decision has major implications for businesses leveraging AI-generated content in marketing, journalism, and creative industries.

Moreover, AI-generated outputs that closely resemble pre-existing copyrighted materials could lead to legal disputes over plagiarism, content originality, and data sourcing ethics (The Role of Generative AI in Managing Industry, 2024). With tools like ChatGPT, Midjourney, and Stable Diffusion enabling users to create text and visual content instantaneously, establishing clear IP frameworks is critical to delineating ownership rights and liability.

To navigate these challenges, companies must develop AI content usage policies that define acceptable AI-generated contributions while ensuring compliance with existing copyright laws. Additionally, collaboration between regulatory bodies, legal experts, and AI researchers is needed to create comprehensive legislation that balances innovation with ethical data usage.

As AI continues to evolve, businesses, governments, and researchers must collaborate to shape policies that align AI innovation with human-centric values. Only through proactive regulation, ethical oversight, and responsible AI integration can the benefits of Generative AI be harnessed while mitigating its most pressing risks.

To ensure responsible AI adoption, organizations must prioritize ethical AI governance, invest in workforce reskilling, and develop clear legal frameworks that protect IP rights and uphold fairness in AI decision-making. The establishment of international AI regulations will be critical in fostering trust, transparency, and accountability in AI deployment.

3.4 Strategic Roadmap for GenAI Adoption

The integration of Generative AI into industrial, commercial, and service-based sectors requires a structured, strategic, and ethical approach. Given the profound economic and operational implications of AI deployment, businesses must ensure that AI adoption aligns with organizational goals, regulatory frameworks, and ethical considerations (GenAI in Business and Industry, 2024). A balanced approach that prioritizes technological innovation while mitigating risks is key to responsible AI integration (GenAI in Business and Industry, 2024).

One of the first steps toward responsible GenAI integration is achieving AI readiness at the organizational level. Companies need to prepare by assessing their technological infrastructure, defining AI objectives, and identifying key stakeholders responsible for AI governance (GenAI in Business and Industry, 2024). As AI reshapes industries, a well-structured change management framework is required to ensure seamless adoption (GenAI in Business and Industry, 2024).

An essential component of AI governance is ensuring algorithmic fairness and reducing bias in AI-generated outputs. AI-driven decision-making systems must be continuously audited to detect and mitigate biases in training data and model predictions (Chui, M., Hazan, E., et al., 2023). Organizations must also foster transparency in AI processes, allowing employees and stakeholders to understand how AI-driven decisions are made (GenAI in Business and Industry, 2024).

Another critical element in responsible AI adoption is proactive regulatory compliance. AI systems must adhere to data protection laws, industry-specific AI standards, and evolving global AI regulations (GenAI in Business and Industry, 2024). Establishing internal AI policies that define acceptable AI applications, ethical AI guidelines, and clear data governance measures is imperative to prevent legal challenges (GenAI in Business and Industry, 2024).

Moreover, the long-term success of AI implementation hinges on continuous AI monitoring and risk assessment. Businesses must establish dedicated AI oversight committees to evaluate AI performance, assess ethical risks, and ensure alignment with corporate strategy (GenAI in Business and Industry, 2024). These governance structures enable organizations to identify potential risks early and implement corrective actions before AI-driven systems cause unintended consequences.

3.4.1 Workforce Upskilling and Ethical AI Frameworks

A well-defined strategic roadmap is essential for businesses seeking to integrate GenAI responsibly while maximizing its benefits and minimizing associated risks. Organizations must prioritize structured AI implementation, balancing technological adoption with ethical AI governance and compliance considerations.

The integration of AI-driven solutions necessitates a fundamental shift in workforce skills and competencies. As AI systems automate repetitive tasks and enhance complex decision-making, businesses must invest in employee training, AI literacy programs, and interdisciplinary AI skill development (GenAI in Business and Industry, 2024).

A GenAI-ready workforce requires employees who possess a blend of technical expertise and adaptive thinking. While traditional roles will evolve, employees must develop proficiency in data analytics, AI model interpretation, and AI-human collaboration techniques (GenAI in Business and Industry, 2024).

Upskilling initiatives should not only focus on technical skills but also emphasize ethical AI considerations, critical thinking, and responsible AI use.

Furthermore, organizations must establish formal ethical AI guidelines to ensure AI applications remain transparent, accountable, and explainable (GenAI in Business and Industry, 2024). Ethical AI governance should emphasize fairness, non-discrimination, and privacy protections, ensuring that AI-generated content and decisions align with human values (GenAI in Business and Industry, 2024).

In addition to ethical guidelines, organizations should implement AI-driven decision-auditing mechanisms. Businesses must be able to trace AI-generated recommendations, assess their impact, and intervene when necessary (GenAI in Business and Industry, 2024). This approach reinforces trust in AI adoption while preventing black-box decision-making that lacks human oversight.

Data governance plays an equally pivotal role in ensuring AI-driven solutions are ethically and legally compliant. Businesses must establish clear policies on AI data handling, informed consent, and secure data storage (GenAI in Business and Industry, 2024). AI applications that interact with sensitive user data—such as healthcare AI and financial AI tools—must undergo rigorous compliance checks to align with regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

Finally, fostering a culture of AI collaboration within organizations is essential. AI should be positioned as an augmentation tool rather than a replacement for human intelligence (GenAI in Business and Industry, 2024). By encouraging cross-functional teams to work alongside AI tools, organizations can enhance innovation, problem-solving, and productivity while minimizing fears of AI-induced workforce disruptions.

3.4.2 Long-term AI sustainability

As AI technologies continue to advance at an accelerated pace, organizations must adopt strategies to future-proof themselves against potential AI disruptions. A sustainable AI strategy involves continuous adaptation, proactive governance, and resilience planning (The Transformative Potential of Generative AI, 2024).

One of the primary strategies for long-term AI sustainability is maintaining corporate agility. Businesses that establish flexible AI governance structures, invest in employee AI education, and continuously refine AI deployment strategies will maintain a competitive edge (The Transformative Potential of Generative AI, 2024). AI readiness should be viewed as an ongoing process rather than a one-time investment, ensuring that organizations can adapt to evolving AI capabilities without facing operational disruptions.

Another key strategy is ensuring that AI deployment aligns with long-term corporate objectives. Companies must assess whether AI-driven innovations contribute to strategic growth, efficiency improvements, and value creation (The Transformative Potential of Generative AI, 2024). Rather than adopting AI for the sake of automation, organizations must ensure that AI addresses real business challenges, enhances decision-making, and fosters sustainable industry progress.

Furthermore, the future of AI-driven businesses will depend on their ability to align AI strategies with environmental, social, and governance (ESG) goals. Sustainable AI deployment should balance economic efficiency with environmental and ethical considerations, ensuring that AI innovation remains responsible and aligned with broader corporate sustainability strategies (The Transformative Potential of Generative AI, 2024).

Finally, businesses must prepare for the next wave of AI evolution by fostering AI research collaborations, industry partnerships, and regulatory engagements. Companies that engage with policymakers, AI researchers, and interdisciplinary AI ethics councils will have a deeper understanding of emerging AI challenges and regulatory expectations, allowing them to anticipate and mitigate potential risks before they escalate.

Chapter 4

Conclusion

4. Conclusion

4.1 Summary of Key Insights

The exploration of Generative AI in various industrial and commercial domains has underscored its transformative potential, while simultaneously revealing the challenges associated with its widespread adoption. From automation and decision-making enhancements to workforce evolution and regulatory challenges, the impact of GenAI is multifaceted, demanding a comprehensive and interdisciplinary approach to its integration (The Transformative Potential of Generative AI, 2024).

The broad spectrum of GenAI's influence spans software engineering, healthcare, business management, and academia, demonstrating its capacity to augment creativity, optimize operations, and drive productivity (The Transformative Potential of Generative AI, 2024). However, despite its economic and strategic benefits, the deployment of GenAI necessitates careful alignment with regulatory, ethical, and organizational frameworks. The necessity for compliance with legal and regulatory requirements, including copyright, data protection, and AI accountability mechanisms, remains paramount to ensure responsible AI use (The Transformative Potential of Generative AI, 2024).

In particular, industries integrating GenAI must balance technological acceleration with ethical governance. Without structured policies and oversight mechanisms, the risk of algorithmic biases, misinformation, and intellectual property conflicts may undermine the credibility and trustworthiness of AI-driven processes (The Transformative Potential of Generative AI, 2024). Organizations must also develop internal AI governance structures that foster transparency, ensuring that AI decision-making remains interpretable and auditable.

As AI-driven automation continues to redefine business operations, the need for workforce transformation is inevitable. The emergence of AI-enhanced workflows is expected to displace certain job functions while simultaneously creating new AI-centric roles. Thus, comprehensive upskilling initiatives must be prioritized to equip employees with the necessary AI literacy, problem-solving, and interdisciplinary collaboration skills. The transition to an AI-augmented workforce will require leadership engagement, strategic workforce planning, and adaptive learning models to facilitate a seamless integration process.

4.2 Further Research on Responsible GenAI Deployment

The unprecedented pace at which GenAI is being deployed has outpaced existing regulatory frameworks, underscoring the urgency for evolving policies that govern AI usage (The Transformative Potential of Generative AI, 2024). Policymakers and industry leaders must work collaboratively to establish ethical AI deployment guidelines, addressing key concerns related to bias mitigation, data privacy, and transparency in AI-driven decision-making.

One of the central elements of responsible AI adoption is the integration of Explainable AI (XAI) principles, ensuring that AI-generated outputs are not only reliable but also interpretable by end-users (Generative AI in UX Design and Research, 2024). The ability to audit and trace AI decisions is fundamental to fostering accountability, particularly in high-stakes applications such as financial services, legal advisory, and healthcare diagnostics. Without mechanisms to validate AI-generated content, organizations risk deploying opaque AI models that may reinforce existing biases or generate misleading outputs.

Further research is also required to assess industry-specific AI practices, ensuring that AI adoption aligns with sectoral regulations and business objectives (Generative AI in UX Design and Research, 2024). As AI-driven UX tools, content generation platforms, and automated decision-making systems become more prevalent, research must explore the evolving role of AI in professional environments. Understanding how UX researchers, business leaders, and policymakers engage with AI tools will provide deeper insights into the practical implementation and long-term effects of AI integration.

Beyond regulatory considerations, further research is needed to explore the psychological and societal impact of AI adoption, particularly regarding human-AI collaboration, employment shifts, and ethical concerns. Addressing how workers interact with AI tools and the potential psychological effects of AI-driven automation will be key areas of exploration. By bridging technical advancements with human-centric research, future studies can ensure that GenAI enhances—rather than disrupts—human productivity and well-being.

4.3 Final Thoughts on Industry-Wide Transformation

As industries navigate the rapid evolution of AI-driven technologies, balancing innovation with responsibility is paramount. While GenAI unlocks unparalleled opportunities for operational efficiency, cost reduction, and creative augmentation, its integration must be approached with a long-term vision for sustainability and ethical governance (GenAI in Business and Industry, 2024).

The evolution from Industry 4.0 to Industry 5.0 is emblematic of this transformation, shifting the focus from automation-centric production models to human-centered, resilient, and adaptive AI applications (The Role of Generative AI in Managing Industry Projects, 2024). In this transition, business leaders, policymakers, and AI researchers must collaborate to ensure that AI deployments reflect not only economic and operational objectives but also ethical and societal considerations.

Looking ahead, organizations must prioritize AI governance structures that ensure accountability, fairness, and security. Establishing cross-disciplinary AI ethics committees can help organizations mitigate risks related to algorithmic biases, data security vulnerabilities, and legal compliance challenges (The Transformative Potential of Generative AI, 2024). Proactively engaging with regulatory bodies and AI standardization initiatives will further strengthen the credibility and sustainability of AI-driven innovations.

Ultimately, the trajectory of AI adoption will be defined by the choices organizations make today. Companies that proactively invest in responsible AI governance, ethical frameworks, and AI literacy programs will be better positioned to thrive in an AI-driven economy. Meanwhile, those that fail to address regulatory concerns, workforce adaptation, and transparency in AI models risk alienating stakeholders and facing legal or reputational challenges.

Thus, as GenAI reshapes industries and professional landscapes, a human-centric approach to AI deployment must remain at the forefront of industry-wide discussions. By aligning technological advancements with ethical safeguards, regulatory compliance, and workforce development strategies, businesses can harness the full potential of AI while ensuring that its adoption remains sustainable, equitable, and beneficial for all stakeholders.

In conclusion, Generative AI represents one of the most transformative technological advancements of the 21st century, with the power to redefine industries, reshape economies, and enhance human creativity. However, its success will be determined not solely by its capabilities but by how effectively it is governed, deployed, and aligned with broader societal values. The future of AI-driven industries will depend on the collective efforts of businesses, researchers, policymakers, and the global workforce to shape AI in a way that is innovative, ethical, and future-ready.

References & bibliography:

- Saha, B. (2024). Generative Artificial Intelligence for Industry: Opportunities, Challenges, and Impact. [Conference paper]. Hitachi Energy, Bangalore, India.
- Cano-Marin, E. (2024). The Transformative Potential of Generative Artificial Intelligence in Business: A Text Mining Analysis on Innovation Data Sources. ESIC Market Economics and Business Journal, 55(2), e333. https://doi.org/10.7200/esicm.55.333
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey & Company, June 2023.
- Mohammed, M. Y., & Skibniewski, M. J. (2023). The Role of Generative AI in Managing Industry Projects: Transforming Industry 4.0 Into Industry 5.0 Driven Economy. Law and Business, 3, 27–41.
- Şahin, O., & Karayel, D. (2024). Generative Artificial Intelligence in Business and Industry: A Systematic Review on the Threshold of Transformation. Journal of Smart Systems Research, 5(2), 156–175. https://doi.org/10.58769/joinssr.1597110
- Kanbach, D. K., Heiduk, L., Blueher, G., Schreiter, M., & Lahmann, A. (2024). The GenAI is out of the bottle: Generative Artificial Intelligence from a Business Model Innovation Perspective. Review of Managerial Science, 18, 1189–1220. https://doi.org/10.1007/s11846-023-00696-z
- Takafoli, M., Li, S., & Mäkelä, V. (2024). Gener`ative AI in User Experience Design and Research: How Do UX Practitioners, Teams, and Companies Use GenAI in Industry? In Proceedings of the Designing Interactive Systems Conference (DIS ’24). ACM. https://doi.org/10.1145/3643834.3660720

- Wolf, V., & Maier, C. (2024). ChatGPT usage in everyday life: A motivation-theoretic mixed-methods study. International Journal of Information Management, 79, 102821.
- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language Models are Few-Shot Learners. In: Advances in Neural Information Processing Systems [Internet]. Curran Associates, Inc.; 2020 [citato 11 aprile 2024]. p. 1877–901. Disponibile su: https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html
- Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding [Internet]. arXiv; 2019 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/1810.04805
- Radford A, Narasimhan K, Salimans T, Sutskever I. Improving Language Understanding by Generative Pre-Training.
- OpenAI, Achiam J, Adler S, Agarwal S, Ahmad L, Akkaya I, et al. GPT-4 Technical Report [Internet]. arXiv; 2024 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2303.08774
- Gemini Team, Anil R, Borgeaud S, Alayrac JB, Yu J, Soricut R, et al. Gemini: A Family of Highly Capable Multimodal Models [Internet]. arXiv; 2024 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2312.11805
- Reid M, Savinov N, Teplyashin D, Lepikhin D, Lillicrap T, Alayrac J baptiste, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context [Internet]. arXiv; 2024 [visitato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2403.05530
- Context by Cohere [Internet]. 2024 [citato 11 aprile 2024]. Command R: RAG at Production Scale. Disponibile su: https://txt.cohere.com/Command-R
- Amazon Web Services, Inc. [Internet]. [citato 11 aprile 2024]. Modelli di fondazione per l'IA generativa – Amazon Titan – AWS. Disponibile su: https://aws.amazon.com/it/bedrock/titan/
- Announcing Jurassic-2 and Task-Specific APIs [Internet]. [citato 11 aprile 2024]. Disponibile su: https://www.ai21.com/blog/introducing-j2

- Simplifying Our Jurassic-2 Offering [Internet]. [citato 11 aprile 2024]. Disponibile su: https://www.ai21.com/blog/simplifying-our-jurassic-2-offering
- Touvron H, Martin L, Stone K, Albert P, Almahairi A, Babaei Y, et al. Llama 2: Open Foundation and Fine-Tuned Chat Models [Internet]. arXiv; 2023 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2307.09288
- Jiang AQ, Sablayrolles A, Mensch A, Bamford C, Chaplot DS, Casas D de las, et al. Mistral 7B [Internet]. arXiv; 2023 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2310.06825
- mistralai/Mistral-7B-Instruct-v0.2 · Hugging Face [Internet]. [citato 11 aprile 2024]. Disponibile su: https://huggingface.co/mistralai/Mistral-7B-In`struct-v0.2
- Gemma Team, Mesnard T, Hardin C, Dadashi R, Bhupatiraju S, Pathak S, et al. Gemma: Open Models Based on Gemini Research and Technology [Internet]. arXiv; 2024 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2403.08295
- google/gemma-7b · Hugging Face [Internet]. 2024 [citato 11 aprile 2024]. Disponibile su: https://huggingface.co/google/gemma-7b
- Hendrycks D, Burns C, Basart S, Zou A, Mazeika M, Song D, et al. Measuring Massive Multitask Language Understanding [Internet]. arXiv; 2021 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2009.03300
- Rein D, Hou BL, Stickland AC, Petty J, Pang RY, Dirani J, et al. GPQA: A Graduate-Level Google-Proof Q&A Benchmark [Internet]. arXiv; 2023 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/2311.12022
- Dua D, Wang Y, Dasigi P, Stanovsky G, Singh S, Gardner M. DROP: A Reading Comprehension Benchmark Requiring Discrete Reasoning Over Paragraphs [Internet]. arXiv; 2019 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/1903.00161
- Zellers R, Holtzman A, Bisk Y, Farhadi A, Choi Y. HellaSwag: Can a Machine Really Finish Your Sentence? [Internet]. arXiv; 2019 [citato 11 aprile 2024]. Disponibile su: http://arxiv.org/abs/1905.07830