Generative AI as Both Subject and Method: Exploring the Future of Work

A Thesis Presented to:

The Honors Tutorial College

Ohio University

In Partial Fulfillment of the Requirements for Graduation from the

Honors Tutorial College

With the degree of

Bachelor of Business Administration

by

Charlie J. Strohm

April 2025

This thesis has been approved by

Dr. Vic Matta

Professor, Analytics & Information Systems

Thesis Adviser

Dr. Raymond Frost

Director of Studies, Business Administration

Dr. Kristina Bross

Dean, Honors Tutorial College

Table of Contents

| Acknowledgements | 5 |
|--|----|
| Abstract | |
| 1 Introduction | |
| 2 Literature Review | |
| 2.1 Generative AI – Definitions and Historical Context | 9 |
| 2.2 Generative AI in Today's Organizations: Inconsistencies Abound | 12 |
| 2.4 Framework Development | |
| 2.4.A. Literature Inclusion | |
| 2.4.B. Generative AI Use for Knowledge Extraction | |
| 3 Methodology | |
| 3.1 Methods for Article Selection | |
| 3.1.A. Academic Literature | |
| 3.1.B. Non-academic Literature | |
| 3.2 Custom GPT | |
| 3.2.A. Custom GPT Training Process | |
| 3.2.B. Custom GPT Comparison Study | |
| 3.2.C. GPT Performance Metrics | |
| 3.3 Knowledge Extraction Process | |
| 4 Results | |
| 4.1 Custom GPT Training Process and Comparison Study Results | |
| 4.2 Generative AI and the Future of Work Analysis Results | |
| 5 Discussion | |
| 5.1 GAI-assisted Knowledge Extraction Process | |
| 5.2 Future of Work Analysis | 40 |
| 5.2.A. Automation and Augmentation | |
| 5.2.B. Hallucination | |
| 5.3.C. Bias | |
| 5.2.D. Ethics | |
| 5.2.E. Regulatory | |
| | |
| 6 Conclusion | |
| Implications | 47 |

| Limitations and Future Research | 48 |
|---|------|
| References | . 50 |
| Appendix 1: Cyclical Training Process Example with Dialog | . 65 |
| Appendix 2: Statement Regarding Generative AI Use | . 69 |

Acknowledgements

As my time at Ohio University comes to a close, I would like to briefly express gratitude to the many figures who have profoundly impacted my professional and personal growth. The Honors Tutorial College provides a unique and enriching educational experience, and I will be forever grateful for the time I spent and the memories I made in this remarkable environment.

Dr. Matta, thank you for your continuous support through the development of this thesis. Your ongoing motivation and expertise were pivotal throughout this process. From weekly meetings to impromptu Sunday morning Teams calls, your commitment to my success is nothing short of amazing. I admire your passion for student success and am grateful for all your support.

Dr. Frost, thank you for your constant support and guidance through the past four years. From our initial HTC seminar, you pushed us to explore our individual interests and passions. Your reassurance to follow my passions was instrumental in the creation of this thesis. I am grateful for your valuable advice during my time as an undergraduate.

To my friends, classmates, and fellow cohort members, thank you for the memories made and constant support these past four years. I will forever be grateful for my time at Ohio University.

Abstract

This study presents an innovative methodology for literature identification and knowledge extraction using Generative Artificial Intelligence (GAI) while addressing key challenges such as hallucinations, selection bias, transparency, and reproducibility. The analysis is focused on the Future of Work as shaped by GAI, a widely discussed topic with recurring themes including automation, augmentation, hallucination, bias, ethics, regulation, and dystopian futures. The proposed methodology systematically examines academic and industry publications to include contemporary perspectives while limiting selection bias. A training process for a custom GPT using natural-language reinforcement training for identifying and extracting key topics is presented, providing a macro-level perspective from a large-scale literature review. The outlined process improves the harmonic mean of precision and recall by over 50% with three iterations of reinforcement learning, identifying the no-code process as effective. This approach is applicable for activities such as market research, financial analysis, political analysis, and public opinion analysis.

1 Introduction

The rapid rise of generative artificial intelligence (GAI) is set to transform cognitive, creative, and organizational domains (Chui et al., 2023). Unlike artificial intelligence of the past, GAI enables the creation of new content at scale and speed. The release of tools such as ChatGPT and DALL-E has brought GAI into public and enterprise consciousness, sparking excitement and concern (Marr, 2023). As organizations race to implement GAI, the academic and professional discourse has become increasingly incongruent, preventing a comprehensive understanding of how GAI will impact the Future of Work (Kinder et al., 2024). This thesis is motivated by the need to systematically identify, analyze, and synthesize the diverse perspectives of GAI's impact on the Future of Work while promoting transparency and replicability, and addressing the lag between technological advancement and academic publishing.

To address these challenges, this thesis presents a novel research-backed methodology that positions GAI as an effective research tool. Specifically, this thesis introduces a structured human-in-the-loop methodology for training and applying an OpenAI custom GPT for literature analysis. This thesis also presents a systematic process for literature identification, balancing credibility and contemporary relevance through GAI-assisted triangulation and traditional database search. Substantively, this thesis provides a macro-level view of the current discourse on GAI and the Future of Work through the extraction of themes and in-depth discussion of findings. By positioning GAI as both the object of study and methodology, this thesis introduces a novel framework for future qualitative research while providing high-level insights on GAI and the Future of Work. The remainder of this thesis is organized as follows. First, a literature review will provide historical context and discuss current perspectives on GAI's organizational impact. The literature review also defines key topics and includes literature regarding the research framework. Next, the research-backed methodology for literature identification, GPT training, and knowledge extraction is presented. Next, the GPT comparison and Future of Work analysis results are provided. Key findings are discussed, exploring GAI as both a research tool and a transformative force for the Future of Work. Finally, a conclusion highlighting the contributions of this thesis, practical implications, and directions for future research is provided.

2 Literature Review

The emergence of Generative Artificial Intelligence (GAI), such as large language models, image generators, and other content-producing systems, represents a turning point in transforming how work will be conducted (Kinder et al., 2024). Unlike earlier forms of artificial intelligence, GAI introduces capabilities that span creative, cognitive, and communicative domains. The rapid advancements of GAI technologies have sparked debate among academics, business leaders, policymakers, and technologists. In short, the diversity of perspectives regarding how GAI will impact the Future of Work prevents a cohesive understanding of the topic (Kinder et al., 2024). This thesis responds to this current problem by presenting a framework for acquiring a macro-level perspective of the topic through systematic, bias-free literature selection and an efficient knowledge extraction process utilizing GAI.

This literature review will provide an overview of generative AI, including a brief historical analysis and examination of recent developments. This is followed by a discussion of key topics within the discourse. Lastly, the conceptual design of the methodology, including literature regarding GAI use in the literature review and applicable training methods, is discussed.

2.1 Generative AI – Definitions and Historical Context

Generative AI is "a system of algorithms or computer processes that can create novel output in text, images, or other media based on user prompts" (NNLM, 2025). This differs from traditional AI, which analyzes and interprets existing data within predefined boundaries, such as predictive analytics or autonomous systems. As the name suggests, generative AI is a subset of AI that can generate information based on existing data (Illinois.edu, 2024).

Generative AI History

Despite common belief, GAI is not necessarily a new technology. Joseph Weizenbaum created ELIZA, a text chatbot that mimicked the work of a psychotherapist with the ability to communicate in natural language in 1966. ELIZA recognized keywords and generated preprogrammed responses, interpreting words as character data. While primitive, ELIZA "opened the path for subsequent advances in the Natural Language Processing domain over the following decades" (Taloka.ai, 2023).

| Welc | ome to | | | | | | |
|---------------------|--|-----------------------------------|-----------------------------|------------------------------|------|--------|-------|
| | EEEEEE | LL | IIII | 2222222 | AA | AAA | |
| | EE | LL | II | 22 | AA | AA | |
| | EEEEE | LL | II | 222 | AAA | AAAA | |
| | EE | LL | II | 22 | AA | AA | |
| | EEEEEE | LLLLLL | IIII | 2222222 | AA | AA | |
| Eliz The This | a is a mock Rogerian original program was implementation by No | psychoth describe orbert La | erapis d by J indstei | t. oseph Wei ner 2005. | zenb | aum in | 1966. |
| ELIZA: | Please tell me what' | s been b | otheri | ng you. | | | |
| YOU: | I don't really know | | | | | | |
| ELIZA: | Do you wish to be ab | le to re | ally k | now ? | | | |
| YOU: | Hmm well, I guess | yes | | | | | |
| ELIZA: | Do you say you guess | yes for | some | special r | easo | | |
| YOU: | I don't really know | | | | | | |
| ELIZA: | Does that trouble yo | ou ? | | | | | |
| YOU: | What do you think? | | | | | | |
| ELIZA: | What are your feelin | igs now ? | | | | | |
| YOU: | Sad | | | | | | |
| ELIZA: | Does talking about t | his both | er you | | | | |
| YOU: | | | | | | | |

Figure 1: A Conversation with ELIZA (Tzouganatou, 2018)

Decades of technological advancements led to the creation of generative adversarial networks (GANs) in 2014. GANs are able to generate accurate images using a generator and a discriminator. The generator generates data, and the discriminator penalizes the generator for producing bad data (Google, 2025). One year later, diffusion models were created, which now power image-generation models such as DALL-E, Midjourney, and Stable Diffusion. Diffusion models work by "destroying training data through the successive addition of Gaussian noise, and then learning to recover the data by reversing the noising process" (O'Connor, 2022).

A significant development in generative AI is the creation of Generative Pre-trained Transformer (GPT) models by OpenAI in 2018. GPT models analyze an input sequence and apply complex algorithms to predict the most likely output. In other words, GPTs "use probability to identify the best possible next word in a sentence based on all previous words" (Belcic & Stryker, 2024). OpenAI released ChatGPT for public use in 2022, allowing users to interact with the application. As noted by Marr (2023), this "marked a pivotal moment when the world started acknowledging this groundbreaking technology...it became clear just how impactful this technology would become".

Since the release of ChatGPT, massive leaps in abilities have been seen as the technology improves. Driving forces behind this improvement can be linked to increases in training compute and hardware improvements. The training compute used to train notable AI models has grown roughly 4.6x yearly (Rahman, 2024). Leopold Ashcenbrenner, an OpenAI researcher, states that this increase in training compute is not the only factor that drives GAI advancement. He believes increases in algorithmic efficiencies, or "effective compute", and "unhobbling" gains (removing

limitations) will advance this technology in the near future. Standardized tests can be used as benchmarks to understand the increase in abilities. Figure 2 presents the performance of GPT-4 (2023) and GPT-3.5 on multiple standardized tests, showing the improvements seen in one year.

| Performance on common exams (percentile compared to human test-takers) | | | | | | |
|---|--|--|--|--|--|--|
| GPT-4 (2023) | GPT-3.5 (2022) | | | | | |
| 90th | 10th | | | | | |
| 88th | 40th | | | | | |
| 97th | 87th | | | | | |
| 99th | 63rd | | | | | |
| 80th | 25th | | | | | |
| 99th | 32nd | | | | | |
| 51st | 3rd | | | | | |
| 80th | 34th | | | | | |
| 92nd | 40th | | | | | |
| 92nd | 51st | | | | | |
| | human test- GPT-4 (2023) 90th 88th 97th 99th 80th 99th 51st 80th 92nd 92nd | | | | | |

Figure 2: Common Exam Performance (2022-2023) (Aschenbrenner, 2024)

Beyond performance increases, advances in tool integration, such as live web browsing, code interpretation, image analysis, and in-platform customization interfaces, have improved these systems' usefulness and potential business use-cases. The term multimodal GAI refers to GAI models that can process multiple data types, including text, images, audio, or other input forms, increasing usefulness and enhancing human-computer interaction (Stryker, 2024). The ability of a model to combine different forms of inputs has been shown to benefit accuracy, contextualization, and problem solving, further advancing GAI capabilities (Curtis & kidd, 2024). To capitalize on this technology, academics and business leaders are asking how GAI will impact the future of work.

2.2 Generative AI in Today's Organizations: Inconsistencies Abound

Contradictions appear in the literature regarding how GAI adoption will impact business. A job can be described as a collection of tasks of varying complexity and importance, and preliminary research reveals that an estimated 80% of workers could see 10% of tasks affected, and 19% of workers could see over 50% of functions affected (Eloundou et al., 2023). The contradictions in the literature arise when attempts are made to predict scenarios of job displacement and/or creation resulting from GAI. Eloundou et al. (2023) state that entirely new professions will arise, while industry figure Mustafa Suleyman argues that professions created by GAI could be performed with GAI (Suleyman, 2023).

This contradiction in how labor markets may be affected by GAI can be traced to the topics of automation and augmentation. Traditional beliefs are that technology can only automate routine tasks such as physical labor, or information collection and processing (Muro et al., 2019), showing a continued need for the innate abilities of the human mind (Jarrahi, 2018). However, with GAI advancement, this range of tasks could include creative, cognitive, and complex tasks (Eloundou et al., 2023; Shact et al., 2024; Thornhill, 2024).

Further incongruities emerge regarding concerns around reliability and transparency. It is understood that GAI systems tend to "hallucinate", or produce outputs that appear correct but are inaccurate and potentially misleading (Mittal et al., 2024; Zysman & Nitzberg, 2024). Additionally, GAI is known to hold potentially harmful biases due to biased training data (Shact et al., 2024). Research shows that business leaders understand these limitations, with 72% of surveyed business leaders stating they worry about the accuracy of GAI outputs (MIT Insights, 2024). While business leaders understand these risks, survey data shows that around 30% of

organizations review only a few generated outputs (Singla et al., 2025). Concerns around GAI risks of hallucination and bias were identified as key topics of discussion in the literature.

As academics, business leaders, and policymakers struggle to find cohesion in how GAI will impact the Future of Work, concerns regarding ethics and regulations emerge in the literature. Legal and ethical issues, including privacy, data protection, and intellectual property, arise with the implementation of GAI (Aanestad, 2024; Woodall, 2024). Literature suggests that rapid technological advancement and a traditionally slower policy development process create a lag in regulation (Shact et al., 2024; Zorpette, 2024). Rapid advancement without corresponding regulatory safeguards could amplify ethical and legal risks, potentially resulting in dystopian scenarios (The Week, 2024). Current AI systems are classified as artificial narrow intelligence, designed to perform specific tasks within limited domains (IBM, 2023). However, emerging concepts in the field anticipate the development of artificial general intelligence and artificial superintelligence. These advanced forms of AI could surpass human cognitive capabilities, introducing significant risks, including the potential for dystopian outcomes (Klingler, 2024; Kuusi & Heinonen, 2022). For these reasons, Ethics, regulatory, and dystopia are identified as key topics of discussion.

Inconsistent opinions presented in the literature motivate the exploration of these topics. In the initial literature review, diverse viewpoints on how GAI will impact the Future of Work were frequent. As discussed above, examination of the literature regarding generative AI and the Future of Work reveals recurring topics. Automation, augmentation, hallucination, bias, ethics, regulatory, and dystopia were identified as key topics of discussion. Each topic will be briefly discussed.

Automation

Automation with GAI refers to delegating tasks traditionally performed by humans to GAI with minimal human intervention. Past technological advances have led to the automation of work in industries such as manufacturing, often due to tasks being restructured to become automated (Frey & Osborne, 2024). The key discussion in the literature surrounding automation is whether GAI will be able to go beyond routine cognitive tasks, automating complex social and creative tasks previously believed to require the nuanced thinking of the human mind.

Augmentation

Augmentation in GAI refers to enhancing human capabilities by implementing GAI technology. Much research examines the benefits of the human-in-the-loop approach to GAI, using this technology as an intelligent assistant (Gruber, 2017; Kuusi & Heinonen, 2022). In brief, augmentation combines the human mind's innate abilities with the strengths of GAI to achieve better outcomes (Jarrahi, 2018).

Hallucination

Hallucination in GAI refers to AI systems generating information that appears credible but is factually inaccurate, misleading, or entirely fabricated (Zysman & Nitzberg, 2024). Hallucination is seen as a hurdle for GAI adoption, and the literature discusses business risks from hallucination.

Bias

Bias in GAI refers to outputs from GAI that are unfair or harmful due to biases in the training data behind the model (Lanamäki et al., 2024; Shact et al., 2024). Bias can result in unfair representations and may reinforce inequalities or present unbalanced perspectives (Fry,

2024). Bias is discussed as a widely known risk of GAI. However, the literature suggests it is most prevalent in image generation applications (Nicoletti & Bass, 2023).

Ethics

The literature emphasizes legal and humanistic ethical concerns that could arise with GAI implementation. Legal ethical themes include personal data protection, privacy concerns, and intellectual property rights (Teli et al., 2024), while humanistic ethical concerns examine potential economic disparities and impacts on employee well-being (Marr, 2024).

Regulatory

The regulatory topic refers to legal regulations regarding GAI. As discussed in the ethics section, privacy and intellectual property concerns are widely discussed. Jurisdictions have begun creating legislation regarding GAI, such as the EU AI Act (Zorpette, 2024). Regulatory discussion is noteworthy as GAI advances rapidly.

Dystopian Futures

Dystopian futures refer to speculative fears regarding the existential risks and catastrophic consequences of AI systems. Although less prevalent, this topic is important as safety and well-being remain paramount with technological development.

2.4 Framework Development

2.4.A. Literature Inclusion

Academic literature is peer-reviewed and reliable, but suffers from a lag in publication (McLean et al., 2023). Contemporary perspectives in online industry articles solve the lag issue for a fast-paced technology, but can suffer from predispositions. If this could be mitigated, then a review of both academic and industry literature may provide a more stable and credible review of the Future of Work with GAI, while addressing the lag issue (Hall, 2017). As noted by Snyder

(2019), "When researchers are selective of the evidence on which to build their research, ignoring research that points the other way, serious problems can be faced". Strategies were employed to combat this bias from the literature selection using a transparent systematic approach for article inclusion (See Methodology 3.1.B.).

2.4.B. Generative AI Use for Knowledge Extraction

Utilizing GAI has been shown to increase the efficiency of literature review (Li et al., 2025), and is a widely discussed topic in academic literature. GAI tools have been shown to streamline systematic literature reviews, reducing time and human errors (Li et al., 2025). Further, research indicates GAI can assist with skimming literature and providing summaries of included articles (Mozelius & Humble, 2024). While using GAI for literature reviews is shown to have benefits, many journals have released guidelines for use in research, including disclosure policies and permission requirements (Hoover, 2023). According to Mozelius & Humble, (2024) using GAI as a support tool is recommended, noting that the main analysis and conclusion should be human-conducted and all its findings should be validated.

To use GAI, it is recommended that you customize, configure, and train a custom GPTbased application with a human-in-the-loop for output validation. This requires fine-tuning, reinforcement learning (using human feedback), and prompt engineering methods.

Fine-tuning

Fine-tuning is adapting a pre-trained model for specific tasks (Bergmann, 2024). In other words, fine-tuning trains the base model on a smaller task-specific dataset, enhancing the model's understanding of a particular concept. Definitions and examples were used to fine-tune the custom GPT.

Reinforcement Learning with Human Feedback

Reinforcement learning with human feedback (RLHF) is a technique that utilizes human feedback to optimize large language models. RLHF incentivizes the model to perform tasks that are more aligned with users' goals (aws, 2025). RLHF was used in the Cyclical GPT Reinforcement Process (see Figure 4).

Prompt Engineering

Prompt Engineering is the "practice of designing inputs for AI tools that will produce optimal outputs" (McKinsey, 2024). Prompt engineering was used throughout the custom GPT training process, with context and specific, clear instructions in all prompts.

This thesis uses a literature-backed approach to develop a custom GPT for knowledge extraction. Training steps outlined in the methodology align with current machine learning training techniques. The human-in-the-loop approach was taken through this process, leveraging AI and human abilities (Jarrahi, 2018). This novel research approach is the main contribution of this thesis.

3 Methodology

The contribution of this paper is twofold. (i) to demonstrate a process of selecting, extracting, and presenting topics in the large-scale literature review using GAI, while addressing key challenges such as hallucinations, selection bias, reliability, transparency, and reproducibility. (ii) to discuss the Future of Work with GAI with clarity and credibility. The methodology for this thesis is outlined below.

- 1. Identify emergent topics in literature (Augmentation, Automation, Hallucination, Bias, Ethics, Regulation, and Dystopia) through initial literature review
- 2. Identify publications (academic/industry) for inclusion
 - a. AIS-endorsed journals searched via applicable keywords and limiters

- b. Identify credible online publishers, and use search tools to identify relevant articles
- 3. Develop a GAI-based custom GPT to parse for topics in the list of articles
 - a. Develop a Process for custom GPT training
 - b. Train and validate custom GPT accuracy through a comparison study
 - c. Extract thematic patterns in the Future of Work with GAI
- 4. Discuss the extracted content

3.1 Methods for Article Selection

The electronic literature search for this paper was divided into two sections: (A) Academic literature and (B) online sources. Published literature from the selected journals is peer-reviewed before publication, so it is inherently more credible than online sources. Papers from online sources do not go through the same peer-review process, so a process for article selection was created to provide credibility, transparency, bias reduction, and replicability. The methods for article selection are as follows.

3.1.A. Academic Literature

The included literature was identified using the Business Source Complete database and the AIS eLibrary. Journals exclusively endorsed by the Association for Information Systems were included. These journals were in the following groups: the AIS Senior Scholars List of Premier Journals, AIS Journals, and AIS-affiliated journals (ais.net.org/page/AIS elibrary). The journals were searched using keywords related to GAI and the future of work. The following Boolean search terms were developed with search appropriate limiters:

TI ('AI' OR 'Artificial Intelligence' OR 'Generative AI' OR 'Generative Artificial Intelligence' OR 'GPT' or 'GenAI' or 'LLM') AND TX ('Transformation' OR 'Future of Work' OR 'Organizational Change' OR 'Workforce Transformation'). Results were limited to articles published in English and published in 2024. The literature search was completed on the 30th of January 2025. Journals included in the academic literature search include:

Senior Scholars list

- Decision Support Systems
- European Journal of Information Systems
- Information & Management
- Information and Organization
- Information Systems Journal
- Information Systems Research
- Journal of the AIS
- Journal of Strategic Information Systems
- MIS Quarterly

AIS Journals

- Journal of the Association for Information Systems
- Communications of the Association of Information Systems
- AIS Transactions on Human-Computer Interaction
- AIS Transactions on Replication Research
- Pacific Asia Journal of the Association for Information Systems
- MIS Quarterly Executive
- Journal of the Midwest Association for Information Systems
- The Journal of the Southern Association for Information Systems
- RELCASI
- Scandinavian Journal of Information Systems

AIS Affiliated Journals

- Business & Information Systems Engineering
- Journal of Information Systems Education
- Journal of Information Systems and Technology Management
- Journal of Information Technology

Journals Excluded

- Sysèmes d'information et Management (published in French)
- Journal of Information Technology Theory and Application (discontinued after 2019)

3.1.B. Non-academic Literature

With non-academic literature, one needs to ensure transparency, replicability, and the minimization of bias (Snyder, 2019). Therefore, a standardized protocol was developed for a balanced selection of publishers from internet search engines (Google and Bing) and prominent GAI applications (ChatGPT and Google Gemini). The following terms were used for each.

Internet Search: "Top online publishers for articles about generative AI"

GAI Prompt: "I want you to act as a researcher at a university. You want to include nonacademic papers in your research and want to use credible online sources. List the top 20 most reliable and well-regarded in-industry publishers for content about generative AI and how it relates to business, future of work, workforce transformation, AI technological advancement, ethics, and regulations."

The process of asking two LLMs and two search engines for reputable online publishers removes convenience bias from the article selection process for online sources. The results from these four sources were recorded, rid of duplicates, and processed through a protocol to increase reliability. Publishers were removed if present in less than two of the four searches. This process is visualized in Figure 3.



Figure 3: Publisher Triangulation Process

Through the above-mentioned process, 15 reputable publishers were identified. A Google site-specific search was conducted to search publications for articles with the keywords 'generative AI' AND 'Future of Work' OR 'Workforce' OR 'Jobs' published after January 1, 2024, using the following query:

Site:publisherURL "Generative AI" ("Future of Work"OR"Workforce"OR"Jobs") after:2024-01-01

Each publisher was individually searched on February 3, 2025. The first three articles based on Google rankings were selected for inclusion in the study. The included publishers are listed below:

- Deloitte
- Financial Times
- Forbes
- Gartner
- Harvard Business Review
- IEEE (Institute of Electrical and Electronics Engineers)

- McKinsey & Company
- MIT Technology Review
- OpenAI
- PwC
- Science Daily
- TechCrunch
- VentureBeat
- The Verge
- Wired

Literature Matrix

This article selection process takes an inclusive approach, allowing a wide range of perspectives. A total of 25 academic papers and 45 online articles were identified as fitting the criteria for this study. The 70 articles provide an in-depth, transparent, and replicable review of the current state of the many perspectives regarding GAI and the Future of Work.

3.2 Custom GPT

This thesis will highlight a GAI utilization process to assist with identifying thematic presence within the included literature. A Custom GPT was built on OpenAI's architecture utilizing the GPT-40 model. OpenAI defines a custom GPT as a "Custom version of ChatGPT that users can tailor for specific tasks or topics by combining instructions, knowledge, and capabilities" (OpenAI, 2024b). The purpose of using GAI is to allow for the efficient identification of thematic presence and the accurate extraction of in-text quotations related to this

research. This study also aims to serve as a testament to the effective utilization of GAI in research by providing insight into training, testing, and validating the accuracy of GAI outputs.

To provide a custom GPT that behaves in a manner suitable for academic research, steps were taken to provide transparency to the training protocol. An important note is that the current state of narrow artificial intelligence and GAI technology is that GPTs have no conceptual understanding of what they are talking about and may hallucinate responses, leading to inaccurate outputs. To validate the idea that a custom GPT could be trained to become a specialized model sufficient for academic studies, a custom GPT was trained using fine-tuning, reinforcement learning from human feedback, and prompt engineering.

3.2.A. Custom GPT Training Process

Reinforcement learning from human feedback (RLHF) is a machine learning technique where human feedback is used to guide the learning process of a model. RLHF leverages human feedback to mitigate issues such as toxicity and hallucinations (Chaudhari et al., 2024). In Step 3 of the comparison study process, the authors use RLHF to train the GPT on three articles. This process is one of the contributions of this thesis, as this research examines how GAI will impact the Future of Work and outlines an effective no-code process for GPT training. The training process is visualized below in Figure 4.



Figure 4: Cyclical GPT Reinforcement Process Diagram

This cyclical process allows the model to summarize the feedback in natural language and use it to reinforce the model further.

Further, uploading the summarized feedback to the knowledge configuration allows the model to refer to it in future analyses. The cyclical nature also allows benchmarked scores to be calculated after a set number of iterations, allowing a more transparent view of the model's performance. Appendix 1 provides a complete example of one iteration, including prompts, responses, feedback, and summary documents.

3.2.B. Custom GPT Comparison Study

The following question was developed to validate the above-described GPT training process. Validating this process produces a GPT capable of extensive, accurate sentiment analysis of literature.

Q1: Can an OpenAI custom GPT be further trained to generate accurate analysis of topical presence in literature?

A comparison study was conducted between an untrained custom GPT and a trained

custom GPT to provide evidence of the above-mentioned question. The process for the

comparison study is outlined below.

- 1. The author reads, identifies, and records instances of topics in four articles.
- 2. Author uploads article four to untrained GPT, prompting with Instructional Prompt 2 (shown below), records results.
- 3. Author opens custom GPT, uses Cyclical GPT Reinforcement Process to train custom GPT on papers 1-3, shown in Figure 4.
- 4. The author uploads article four to custom GPT, prompting with Instructional Prompt 2, records results.
- 5. The author compares output from steps 2 and 4 with personal analysis, providing proof-of-concept

For Replicability purposes, Instructional Prompt 2 is shown below.

Instructional Prompt 2: "You are to act as a researcher at a university researching generative AI and the future of work. You will be given academic papers and online articles. Your job is to read and comprehend these articles. Then you will search these articles for seven topics being (i) Automation, (ii) Augmentation, (iii) Hallucination, (iv) Bias, (v) Regulatory, (vi) Ethics, and (vii) Dystopian futures. You will identify when these topics are present and extract the language from the papers that shows these topics are presented. You will need to provide the exact quote and page number where the retrieved text exists. If you are ever confused or unclear, you will reply with "Unable to understand". If this text does not exist in the document do not generate it, instead state "Does not exist". You will be trained based on a comparison of your analysis and my analysis. For each topic, you need to provide two things. First, a bullet point and a paraphrased explanation of whether or not the topic is present. Next, a second bullet providing the exact quotes from the document for each topic. Each time you do this, make sure that you have evaluated all instances of that topic. Also, list the number of times each topic is discussed within the uploaded text. For some topics you may have zero instances, in which you will reply "Does not exist", and for some topics you may have multiple instances for which you will then provide the exact text for each instance."

The initial results from the untrained analysis provide evidence that an OpenAI custom

GPT can identify the presence of topics and extract exact text. The output from the untrained

GPT was recorded and saved for later comparison. With the untrained output saved, the custom

GPT was trained with three iterations of the cyclical GPT reinforcement process. Article four was uploaded to the trained GPT, and the results were recorded. With both the untrained and trained models' outputs recorded, metrics for calculating performance were developed and described below.

3.2.C. GPT Performance Metrics

To provide a definitive answer to the previously stated question *can an OpenAI custom GPT be further trained to generate accurate analysis of topical presence in literature,* a comparison of the output between the untrained GPT and the custom GPT was conducted. With the comparison study outputs saved, a systematic process for scoring the custom GPT was developed. The established metrics to use in this process include precision, recall, and F1 score, providing a nuanced understanding of the model's accuracy (McDonough, 2024). These metrics and what they represent will be briefly explained.

 $Precision = \frac{True \ Positive}{True \ Positive + False \ Positive} \qquad Recall = \frac{True \ Positive}{True \ Positive + False \ Negative}$

Precision measures the accuracy of the GPTs' topical presence instances. It is calculated as the ratio of true positives to the sum of both true and false positives. Recall, or true positive rate, calculates the model's ability to find all relevant instances in a data set (Urwin, 2024).

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall}$$

The F-1 score is the harmonic mean of precision and recall, providing a single metric to measure a model's performance (Mahedi Hasan et al., 2017). It balances caution (precision) and thoroughness (recall). The F-1 score is valuable since it considers both false positives and false negatives.

The detailed results of this comparison study can be found in the results section of this thesis. The initial results provided a definitive answer to the above-stated *Q1*, providing evidence that an OpenAI custom GPT can be further trained to accurately identify the presence of topics within literature, allowing for increased efficiency in large-scale literature knowledge extraction with a high degree of accuracy. This proof-of-concept validated our process for GAI use in this thesis. The process for applying the trained custom GPT is provided below.

3.3 Knowledge Extraction Process

For this study, a harmonized benchmark F1 score of ≥ 0.90 was deemed acceptable for application. The cyclical GPT reinforcement training process shown in Figure 4 was continued for five iterations, with a resulting F1 Score of 0.93. With the model fully trained, selected articles were uploaded and analyzed in batches of 10 to prevent tokenization issues. The resulting outputs from each batch were exported to Excel for analysis to avoid possible tokenization issues. As noted by Jarrahi (2018), humans and GAI systems working synergistically produce better outcomes. To utilize the "human-in-the-loop" approach, accuracy checks were conducted via human analysis every ten documents using text extraction verification and F1 score calculations. If accuracy inconsistencies were identified, the model was further trained via the cyclical GPT reinforcement protocol to achieve an F1 score ≥ 0.90 . This process is shown below. The complete literature matrix of 45 in-industry publications and 25 academic papers was uploaded to the trained GPT for analysis.



Figure 5: Knowledge Extraction Process Flowchart

Uploading articles in batches of 10 not only reduces the risk of tokenization limits that could hinder accuracy, but it also allows for a deeper understanding of how the model's performance changes after multiple analyses. By calculating performance metrics through the knowledge extraction process, data providing insight for future applications of this process is discovered. The results of the GPT comparison study and the Future of Work literature analysis are found in the next section of this thesis.

4 Results

The results section will examine both the results of the custom GPT comparison study and the Future of Work as two contributions of this thesis.

4.1 Custom GPT Training Process and Comparison Study Results

The GPT training process and comparison study were created to answer the following question:

Q1: Can an OpenAI custom GPT be further trained to generate an accurate analysis of topical presence in literature?

The results of the comparison study show that after three iterations of the cyclical GPT reinforcement training process, notable improvements are seen in the model's ability to both accurately identify the presence of topics within literature and extract the exact text containing the sentiment of the associated topic. Following the process outlined in the methodology, the comparison study was conducted. With the outputs recorded, confusion matrices were developed and performance metrics calculated. Tables 1 and 2 provide confusion matrices, and Figure 4 presents the performance metrics from the trained and untrained models in a comparative visualization.

Untrained GPT

Trained GPT

| | | Act | ual | | | Act | ual |
|---------|-------------|---------|-------------|---------|-------------|---------|-------------|
| | | Present | Not Present | | | Present | Not Present |
| edicted | Present | 7 | 2 | edicted | Present | 14 | 3 |
| GPT Pr | Not Present | 8 | | GPT Pr | Not Present | 1 | |

Tables 1 & 2: GPT Confusion Matrices

With confusion matrices developed, performance metrics were calculated for both models to provide insight into performance and answer the above-stated *Q1*. This proves that a custom GPT can be effectively further trained to analyze topical presence in literature accurately.





The results from this exercise show that three iterations of the Cyclical GPT Reinforcement Process led to increases in all three metrics chosen for performance calculations. Three iterations led to over 50% improvement in the harmonized mean (F-1) of precision and recall, definitively proving that a custom GPT can be further trained for this application. Further, this result validates the Cyclical GPT Reinforcement Process as an effective natural language process for increased model performance.

As outlined in our methodology, accuracy checks were conducted every 10 documents via performance calculations and exact text extraction verification to ensure accuracy through the knowledge extraction process. These performance calculations provide deeper insight into the relationship between the model's analytical processes and the resulting accuracy through the knowledge extraction process. The chronologized F-1 scores are visualized in Figure 7.



Figure 7: F-1 Score Progression

Performance decreased during subsequent article analysis, highlighting the need for human involvement when following this methodology. Analysis of model performance throughout the knowledge extraction process further emphasizes the benefit of the human-in-theloop approach to working with GAI. Calculating performance metrics at predetermined intervals allowed for the quick correction of irregularities using further reinforcement learning.

The results from this comparison study validate the Cyclical GPT Reinforcement Process as an effective no-code process for model performance enhancement. It was determined that an OpenAI custom GPT can identify the presence of topics and extract exact text from literature, and a custom GPT can be further trained using a systematic process to generate more accurate analysis. This process could be applied to tasks that require efficient large-scale review of perspectives, such as, but not limited to, financial analysis, political analysis, market research, or public opinion analysis.

4.2 Generative AI and the Future of Work Analysis Results

For this research, 70 articles were analyzed by a trained OpenAI custom GPT for stringent analysis. Academic articles were identified via keyword search from AIS journals, and industry publications were identified using the Publisher Triangulation Process and a Google site-specific search. A custom GPT was extensively trained using the Cyclical GPT Reinforcement Training Process to accurately identify the presence of the topics of automation, augmentation, hallucination, bias, regulatory matters, ethics, and dystopian possibilities. The results are presented in Table 3 as unprocessed data directly outputted from the trained GPT with academic sources highlighted in grey.

| Citation | Automation | Augmentation | Hallucination | Bias | Regulatory | Ethics | Dystopian Futures |
|------------------------------------|------------|--------------|---------------|------|------------|--------|----------------------|
| (Aanestad, 2024) | 0 | 2 | 0 | 1 | 2 | 1 | 0 |
| (Carroll et al., 2024) | 3 | 3 | 1 | 2 | 1 | 2 | 4 |
| (Constantinides et al., 2024) | 4 | 5 | 2 | 1 | 1 | 1 | 2 |
| (Cui et al., 2024) | 3 | 5 | 0 | 2 | 1 | 3 | 1 |
| (Woodall, 2024) | 2 | 6 | 0 | 0 | 1 | 1 | 1 |
| (University College London, 2024a) | 0 | 2 | 4 | 2 | 0 | 1 | 0 |
| (University College London, 2024b) | 0 | 1 | 2 | 5 | 0 | 1 | 0 |
| (lascaze et al., 12/24) | 4 | 5 | 1 | 2 | 1 | 3 | 0 |
| (Mittal et al., 2024) | 6 | 7 | 2 | 3 | 2 | 4 | 1 |
| (Shact et al., 2024) | 5 | 6 | 3 | 4 | 2 | 5 | 2 |
| (Dos Santos & Williamson, 2024) | 2 | 4 | 0 | 1 | 1 | 0 | 0 |
| (Hamirani, 2024) | 4 | 5 | 0 | 2 | 1 | 3 | 0 |
| (Marr, 2024) | 5 | 6 | 0 | 3 | 2 | 4 | 0 |
| (McKendrick, 2024) | 4 | 5 | 0 | 2 | 2 | 5 | 0 |
| (Financial Times, 2024a) | 3 | 6 | 0 | 3 | 2 | 4 | 0 |
| (Financial Times, 2024b) | 4 | 5 | 1 | 2 | 3 | 5 | 1 |
| (Thornhill, 2024) | 5 | 4 | 0 | 2 | 1 | 3 | 2 |
| (Poitevin, 2024) | 1 | 5 | 0 | 1 | 0 | 0 | 0 |
| (Keen, 2024) | 1 | 3 | 0 | 0 | 0 | 0 | 0 |
| (Goasduff, 2024) | 3 | 5 | 0 | 0 | 0 | 0 | 0 |
| (Hadidi & Klein, 2024) | 2 | 4 | 0 | 1 | 2 | 1 | 0 |

| Citation | Automation | Augmentation | Hallucination | Bias | Regulatory | Ethics | Dystopian Futures |
|-------------------------|------------|--------------|---------------|------|------------|--------|----------------------|
| (Demirci et al., 2024) | 4 | 3 | 0 | 1 | 0 | 0 | 1 |
| (Lovich et al., 2024) | 0 | 5 | 0 | 1 | 0 | 2 | 0 |
| (Rosani & Farri, 2024) | 0 | 3 | 0 | 0 | 0 | 0 | 0 |
| (Holmström, 2024) | 1 | 4 | 0 | 1 | 2 | 2 | 1 |
| (Huy et al., 2024) | 3 | 4 | 1 | 2 | 1 | 2 | 0 |
| (Potter, 2024) | 3 | 5 | 0 | 0 | 0 | 1 | 1 |
| (Pretz, 2025) | 3 | 5 | 0 | 0 | 0 | 1 | 0 |
| (Zorpette, 2024) | 3 | 4 | 0 | 0 | 2 | 1 | 0 |
| (Lanamäki et al., 2024) | 3 | 5 | 1 | 2 | 3 | 2 | 2 |
| (Lang et al., 2024) | 1 | 3 | 1 | 1 | 0 | 2 | 0 |
| (C. Li et al., 2024) | 3 | 4 | 1 | 3 | 2 | 2 | 1 |
| (Livari et al., 2024) | 0 | 4 | 0 | 2 | 0 | 3 | 1 |
| (Ma et al., 2024) | 0 | 2 | 3 | 2 | 1 | 2 | 1 |
| (Maurya et al., 2024) | 2 | 4 | 0 | 2 | 1 | 2 | 1 |
| (De Smet et al., 2024) | 3 | 5 | 0 | 0 | 1 | 2 | 1 |
| (Hazan et al., 2024) | 4 | 5 | 0 | 0 | 1 | 1 | 1 |
| (Mayer et al., 2025) | 3 | 5 | 1 | 1 | 1 | 2 | 1 |
| (MIT Insights, 2024) | 3 | 5 | 1 | 0 | 0 | 1 | 0 |
| (Rotman, 2024b) | 4 | 6 | 1 | 0 | 1 | 2 | 1 |
| (Rotman, 2024a) | 5 | 6 | 0 | 0 | 1 | 1 | 2 |
| (Nyman et al., 2024) | 1 | 4 | 0 | 2 | 1 | 3 | 2 |

| Citation | Automation | Augmentation | Hallucination | Bias | Regulatory | Ethics | Dystopian Futures |
|-------------------------------|------------|--------------|---------------|------|------------|--------|----------------------|
| (OpenAI, 2024a) | 4 | 6 | 0 | 0 | 1 | 1 | 1 |
| (OpenAI, 2024d) | 4 | 6 | 0 | 1 | 1 | 1 | 0 |
| (OpenAI, 2024c) | 3 | 6 | 0 | 0 | 1 | 1 | 0 |
| (PwC, 2024b) | 4 | 6 | 0 | 1 | 1 | 2 | 1 |
| (PwC, 2024a) | 4 | 5 | 0 | 0 | 1 | 1 | 1 |
| (PwC, 2024c) | 4 | 7 | 1 | 1 | 1 | 1 | 1 |
| (Queiroz et al., 2024) | 3 | 3 | 0 | 0 | 1 | 1 | 0 |
| (Saffarizadeh et al., 2024) | 2 | 4 | 0 | 2 | 1 | 3 | 1 |
| (Shahid & Mishra, 2024) | 2 | 4 | 0 | 1 | 2 | 3 | 0 |
| (Stohr et al., 2024) | 4 | 3 | 0 | 1 | 2 | 2 | 0 |
| (Sundberg & Holmström, 2024a) | 3 | 4 | 1 | 2 | 2 | 3 | 1 |
| (Sundberg & Holmström, 2024b) | 2 | 5 | 1 | 1 | 1 | 3 | 0 |
| (Tao et al., 2024) | 3 | 6 | 1 | 2 | 2 | 3 | 1 |
| (Lunden, 2025) | 3 | 5 | 0 | 0 | 1 | 1 | 1 |
| (Miller, 2024) | 2 | 4 | 0 | 0 | 1 | 1 | 1 |
| (Wiggers, 2025) | 0 | 3 | 0 | 0 | 1 | 1 | 1 |
| (Teli et al., 2024) | 4 | 3 | 0 | 2 | 3 | 2 | 1 |
| (Tutun et al., 2024) | 3 | 6 | 1 | 2 | 2 | 4 | 1 |
| (McDermott, 2024) | 2 | 5 | 1 | 1 | 1 | 1 | 0 |
| (Pasch, 2024) | 2 | 6 | 1 | 0 | 0 | 1 | 1 |
| (Plumb, 2024) | 3 | 5 | 0 | 0 | 1 | 1 | 0 |

| Citation | Automation | Augmentation | Hallucination | Bias | Regulatory | Ethics | Dystopian |
|---------------------|------------|--------------|---------------|------|------------|--------|-----------|
| | | | | | | | Futures |
| (Feiner, 2024) | 2 | 3 | 0 | 0 | 2 | 1 | 0 |
| (Larson, 2024) | 3 | 4 | 0 | 0 | 0 | 1 | 1 |
| (Warren, 2024) | 4 | 4 | 1 | 0 | 1 | 1 | 1 |
| (Wang et al., 2024) | 5 | 6 | 0 | 1 | 2 | 3 | 1 |
| (Knight, 2024) | 3 | 4 | 0 | 0 | 1 | 1 | 2 |
| (Marcus, 2024) | 0 | 1 | 3 | 0 | 0 | 1 | 2 |
| (Roslansky, 2024) | 2 | 5 | 0 | 0 | 0 | 1 | 0 |

Table 3: Custom GPT Knowledge Extraction Output

This table presents the unprocessed data from the custom GPT output, which was exported to Excel for further analysis. Through this process, 70 articles were efficiently analyzed, and figures were created, allowing for macro-level topic comprehension in the widerange literature review. Figures were developed to provide a deeper understanding of this unprocessed data. Figure 8 visualizes the total instances identified for each topic.



Figure 8: Total Topic Instances

The results were then proportionately normalized by count to express these figures as proportions relative to the total count, allowing for easier comparison (Funk, 2024) using the formula:

Normalized Proportion =
$$\frac{Topic \ Instance \ count \ (source)}{Article \ Count \ (source)}$$

Figure 9 presents the normalized figures categorized by source type (academic/industry). This eliminates issues arising from unequal sample sizes between article sources and allows for fair comparison.



Figure 9: Normalized Topic Presence

The normalized results can be used to compare the results for academic and online sources. While academic literature and industry reports emphasize Automation and Augmentation, their perspectives diverge. Academic discourse highlights concerns about Ethics, Bias, Hallucinations, and potential Dystopian outcomes. In contrast, leading industry analyses concentrate more on the practical implications of GAI. This focus examines how GAI can streamline processes, boost productivity, and reshape job functions through intelligent automation and augmentation.

5 Discussion

This discussion will be organized into two distinct sections of this paper's contribution: (i) the process of using GAI for knowledge extraction and (ii) the discussion on the Future of Work with GAI.

5.1 GAI-assisted Knowledge Extraction Process

The findings from the custom GPT comparison study emphasize the utility of GAI, specifically OpenAI custom GPTs, in literature analysis. This study demonstrates that a custom GPT can be effectively further trained through natural language reinforcement to accurately identify the presence of predefined topics and extract the associated textual context. Notable improvements were observed in all performance metrics, validating the Cyclical GPT Reinforcement Process as an effective, no-code, and replicable process for increased output accuracy. A 50% improvement in F-1 score was observed through three iterations of the Cyclical GPT Reinforcement Process, further highlighting the practical value of this process.

An important insight gathered through the knowledge extraction process is that model performance decreased during subsequent article analyses, showing a gradual drift in model accuracy without continuous reinforcement. This aligns with literature calling for human oversight with GAI use; the human-in-the-loop (Jarrahi, 2018; Zysman & Nitzberg, 2024). The chronological tracking of F-1 scores and continued feedback allowed this performance drift to be addressed rapidly.

Overall, the results of this process validate GAI use to attain a macro-level perspective from wide-range literature analysis using systematic research-backed training processes. The validation of the Cyclical GPT Reinforcement Process highlights many potential practical applications beyond academic literature review for custom GPT use. Activities that rely on accurate qualitative data analysis, such as financial analysis, market research, and public opinion assessment, could benefit from custom GPTs trained with this process.

5.2 Future of Work Analysis

This thesis aimed to classify emerging perspectives on the impact of generative AI on the Future of Work by developing a systematic framework using AI-driven qualitative analysis. The results revealed strong emphasis on the topics of augmentation, automation, hallucination, bias, regulations, and dystopian futures. The following discussion will follow the outline of first stating widely held opinions, detailing their relevance to the Future of Work, and my key findings and personal perspective for each topic.

5.2.A. Automation and Augmentation

Across academic and industry discourse, automation and augmentation are the most frequently discussed topics regarding GAI's impact. Although separate topics, an interconnectedness exists due to the fact that the automation of routine tasks leads to broader augmentation of workflows, improving overall productivity (Hadidi & Klein, 2024; PwC, 2024b). Due to this interconnectedness, this discussion will examine these topics in tandem while exploring the sentiment present in current publications.

Recent advancements in GAI technology challenge the long-standing notion that only routine and manual tasks are susceptible to automation. As noted by Shact et al. (2024), "Gen AI has taken over software design, code debugging, and customer interactions", domains that require adaptive, cognitive skills. It is noteworthy that customer interaction is present due to human conversation's inherently unpredictable and open-ended nature. Similarly, Thornhill (2024) draws parallels to the industrial revolution, noting "Smart machines will automate brain power in the same way dumb machines automated brawn power during the industrial revolution". These perspectives signal a widespread recognition that previously "safe" jobs are increasingly at risk of GAI automation (Constantinides et al., 2024). Stated plainly, as GAI technology advances, the tasks and jobs that could be impacted increase.

This shift holds implications for the Future of Work. On one hand, concerns around job displacement are prevalent, particularly as GAI gains capability in unstructured environments (Feiner, 2024; Mittal et al., 2024; Potter, 2024). On the other hand, GAI is seen as a tool for worker empowerment. Early and middle-level employees can benefit from increased productivity, faster skill development, and increased relevancy at work, suggesting worker enhancement rather than replacement (Financial Times, 2024a; Knight, 2024; lascaze et al., 12/24; Tao et al., 2024; Thornhill, 2024). An important note is that concerns regarding technological job displacement are not new. According to Lanamäki et al., (2024) this concern has been present for decades, and the latest wave of AI-driven fear is merely a continuation of these fears.

From my analysis, I believe GAI will be a complementary technology in the near term. Rather than job replacement, low-level tasks will be automated, allowing workers to focus on higher-value activities. I believe the human mind's ability to think creatively and connect interpersonally will continue to hold value in a GAI-enhanced enterprise. By embracing this human-centric vision, organizations can capitalize on the value GAI technology provides while enhancing productivity, fostering innovation, and driving growth.

5.2.B. Hallucination

Hallucination, defined broadly as the generation of plausible yet inaccurate or misleading information by GAI models, remains a significant barrier for effective adoption (Lang et al., 2024; Mittal et al., 2024; Warren, 2024). At the technical level, GAI operates as advanced

predictive machines that generate outputs based on token patterns in their training data, an "autocomplete on steroids" (Marcus, 2024). GAI models inherently have no conceptual understanding of what they are talking about, which can result in confident-sounding but factually incorrect statements (Frey & Osborne, 2024; University College London, 2024a). This tendency to fill in the blanks with fabricated information contributes to the hallucination challenge.

The implications of hallucination for the Future of Work are seen as substantial. From a business standpoint, hallucinated outputs pose reputational risks and demand increased oversight (Shact et al., 2024). Survey data reveals over 70% of business leaders have concerns regarding the quality of their AI system's outputs (MIT Insights, 2024), showing this concern is present in high-level decision makers. According to Shact et al. (2024) the risks of GAI hallucination influence how organizations approach GAI, preventing full-scale implementation without human oversight. Hallucination concerns diminish GAI's potential productivity gains and reinforce the need for risk-aware integration. Despite these limitations, organizations are eager to invest and implement GAI, provided these risks can be managed (MIT Insights, 2024).

Based on my analysis, I believe hallucination will remain a key concern in the near future. During this research, I encountered instances of hallucination even after extensively training my custom GPT. I believe the human-in-the-loop approach will continue to be the best practice for GAI use. I believe the potential reputational risks of inaccurate outputs should be enough to prevent organizations from using GAI as a blind-trust automation tool. The current state of GAI should be a complementary tool to enhance workers' abilities. In my view, human oversight remains key with GAI implementation to prevent hallucination.

5.3.C. Bias

Bias in GAI refers to the generation of unfair, harmful, or skewed outputs due to biased training data (Lanamäki et al., 2024; Shact et al., 2024). Frontier models such as GPT-4 are estimated to be trained on roughly 10 trillion words of textual data (Schreiner, 2023). Given this large amount of training data, the inclusion of biased information or a skewed representation of the world in the training data can occur (Gordon, 2023). As earlier noted, models predict what text to generate based on their training data, so humanistic biases present in the training data can be present in outputs (*Science Daily*, 2024; Shact et al., 2024). These biases are especially visible in image-generation models such as Stable Diffusion and Google Gemini, with independent analysis revealing these tools often escalate racial and gender disparities beyond those seen in real-world contexts (Nicoletti & Bass, 2023).

Another form of bias related to GAI is the bias of the user. Emerging research reveals that as humans interact with GAI systems, the model can adopt or amplify their bias (Lanamäki et al., 2024; University College London, 2024b). This could be due to the model deciding to provide an output of interest to the user rather than considering moral issues (Ma et al., 2024). Essentially, if a model is trained to be helpful and agreeable, the model may output what it believes you want to hear, leading to confirmation bias.

From my review, I believe bias is a core limitation that further substantiates the need for human oversight. If GAI systems can not address these issues, tasks that involve subjective judgment or nuanced decision making will not be automated by prediction-based generation. I believe bias will remain a risk in the near future, and the path forward must consist of workforce awareness of this risk. This awareness will allow GAI to be used in applicable

situations as part of augmented workflows, giving technology time to advance and resolve the problem.

5.2.D. Ethics

As GAI technology transforms the enterprise landscape, ethical implementation has emerged as a topic of discussion. An emergent perspective in the literature emphasizes that organizations must comply with legal regulations and ensure alignment with their values and ethical standards (Financial Times, 2024a; Mittal et al., 2024). Further, organizations must develop transparent and trustworthy approaches to GAI use, encouraging the establishment of ethical guidelines and incentivizing ethical use (McKendrick, 2024; Shact et al., 2024). The literature suggests that the ethical use of GAI should not be left to individual interpretation but should instead be defined and managed through organizational strategy.

Ethical considerations related to GAI span both legal and humanistic domains. Legal concerns include data protection, privacy, intellectual property, and compliance with evolving regulations (Aanestad, 2024; Teli et al., 2024; Woodall, 2024). These issues increasingly fall on organizations to address, as they ensure lawful GAI integration. Humanistic ethical concerns around employee well-being are also present in the literature. As noted above, GAI can automate simple tasks, allowing workers to focus on high-value tasks (PwC, 2024b). Literature suggests this shift could lead to increases in employee burnout (Larson, 2024). Further, calls are made to preserve the human aspects of work that promote purpose and emotional welfare (Larson, 2024; Marr, 2024).

From my analysis, the ethical implementation of GAI will be a challenge for organizations in the near future. I believe organizations must go beyond regulatory requirements and create a values-driven culture around GAI use. I believe clear policies that define ethical

boundaries and provide employees with practical guidance regarding GAI use will hold tremendous value. Without these policies, individual interpretations of "ethical use" could vary. I believe aligning legal compliance and organizational values with practical guidance will lead to the ethical GAI implementation within organizations.

5.2.E. Regulatory

Literature calls attention to the misalignment between the pace of technological advancement and the slower evolution of policy (Shact et al., 2024). Governments are beginning to recognize this gap, launching the Workforce of the Future Act to investigate GAI's sectorspecific impact on employment (Feiner, 2024; Knight, 2024). The above-mentioned risks tied to GAI, including data privacy, intellectual property, and ethical use, question whether regulation will keep pace with innovation or if this burden will fall upon organizations to address (Constantinides et al., 2024; Hamirani, 2024; PwC, 2024c).

The regulatory landscape around GAI will influence how organizations implement this technology. Inconsistent or outdated regulations may slow adoption. Examples of this outdated regulation include the E.U. AI Act, which approached AI regulation through a "product safety" blueprint in 2021, assuming stated intended purposes for this technology (Zorpette, 2024). Zorpette (2024) notes this model no longer aligns with multimodal general-purpose AI tools of the present, highlighting the challenges of developing adaptable policies that can evolve alongside the technology it attempts to regulate.

From my analysis, I believe regulations will struggle to keep up with the pace of technological advancement effectively. Examining the E.U. AI Act, four years of advancement rendered the regulation outdated for practical purposes. I believe privacy and data protection regulations likely follow current technology regulations. However, the topic of intellectual

property remains to be seen. The lack of transparency presented by major industry players surrounding training data sources and uncertainty regarding ownership of GAI-produced outputs raises considerable intellectual property concerns. Similar to ethics, providing employees with policies and guidelines and shared responsibility between governments and organizations can foster an environment where innovation and accountability coexist.

5.2.F. Dystopian Futures

A subset of literature presents a dystopian picture of the Future of Work with GAI driven by concerns of unchecked development, job displacement, economic inequality, and the erosion of trust in key institutions. Without sufficient oversight, GAI has the potential to cause harm (Carroll et al., 2024; Mittal et al., 2024). Some perspectives even paint GAI as a threat to democratic discourse through its potential to erode trust in the news (Financial Times, 2024b).

These dystopian concerns could carry tangible implications for how organizations, policymakers, and workers prepare for the future. Research indicates GAI may disproportionately affect low-skilled workers, potentially deepening socioeconomic divides (Wang et al., 2024). Further, AI-driven algorithmic management practices may amplify workplace stress through constant monitoring (Cui et al., 2024).

From my analysis, the dystopian perspectives regarding GAI seem to be a minor, further speculative topic. I believe GAI will advance in the near future, but not to the extreme of existential threats or mass job displacement. My opinion echoes that of the ethics and regulatory topics. I believe adaptive regulation and value-oriented use guidelines can prevent the dystopian futures presented by GAI.

6 Conclusion

This research examined the impact of generative AI on the Future of Work. Through a preliminary literature review, evidence of contradictory perspectives was quickly identified. The ambiguous nature of this topic demanded a novel approach to literature review, capturing and classifying contemporary perspectives efficiently. A research-backed GAI-assisted methodology for article selection and knowledge extraction was developed to identify macro-level insights from academic and industry sources.

Three methodological contributions are presented in this thesis: Literature identification through GAI-assisted triangulation (Figure 3), effective natural language GPT training through Cyclical GPT Reinforcement Process (Figure 4), and replicable GAI-assisted knowledge extraction (Figure 5). This methodology identified a matrix of reputable, relevant, and current literature. A custom GPT was trained to recognize the discussion of relevant topics from literature with >90% accuracy, and these topics were systematically extracted with human oversight. This process demonstrates the utility of GAI use for qualitative research and presents a novel methodology for future applications.

Applying this methodology to the future of work literature, the results highlighted that the current literature suggests a near future of augmented workflows with GAI serving as a complementary technology. Further, transparent and value-driven guidelines for GAI use within organizations, coupled with workforce awareness, are needed to mitigate GAI-associated risks.

Implications

The findings of this thesis hold implications for both research and practice. A methodology for utilizing GAI in qualitative analysis that balances efficiency and accuracy through a structured human-in-the-loop process is established. The validation this process

highlights many potential practical applications beyond academic literature review for custom GPT use. Activities that rely on accurate qualitative data analysis could benefit from custom GPTs trained with this process.

While prior research generally agrees on the transformative nature of GAI, this research provides deeper insight into specific topics of interest in academic and industry discourse. A key insight of this research is the observed difference in emphasis between academic and industry literature. Academic publications are shown to reflect phenomena that have already emerged, focusing heavily on the mitigation of GAI risks such as hallucination and bias. In contrast, industry publications emphasize practical implications driven by market demands and technological advancement. This observed difference in emphasis can be observed in Figure 9, with normalized presence allowing fair comparison.

This thesis contributes to common knowledge by going beyond identifying reoccurring topics in GAI discourse by showing distinctions between academic reflection and industry urgency. Through extensive discussion surrounding these topics, academics, industry leaders, are given a comprehensive understanding of the contemporary and credible discourse surrounding GAI and the Future of Work.

Limitations and Future Research

While this thesis offers meaningful contributions both methodologically and substantively, several limitations should be acknowledged to guide future research. Despite extensive training through natural language reinforcement, the custom GPT still had the limitation of potential hallucination. The 90% model accuracy figure was deemed acceptable for this thesis, and I took steps to validate outputs further. The constraints posed by hallucination could limit the usefulness of the developed methodological processes.

While this thesis adopted an inclusive article selection process, only articles published in English between 2024 and 2025 were included. Further, online sources were identified through search engine ranking, potentially excluding insightful perspectives. Additionally, this thesis focused on predefined topics based on the initial literature review. This focused the analysis and allowed for a structured process, but emergent or niche topics not captured in the initial literature review may have been excluded. Topic rigidity could inadvertently constrain the richness of interpretation in this rapidly evolving research domain.

This thesis utilized an OpenAI GPT-4o custom GPT. Since the completion of this research, newer models have been released, potentially further improving capabilities. Future research could explore the accuracy and trainability of different models, such as those offered by Anthropic, Google, and Meta. Additionally, future research could conduct a longitudinal analysis of GAI discourse over multiple years. This could provide insight into how perspectives shift over time in response to technological improvement.

References

- Aanestad, M. (2024). Workplace Datafication that Serves the Workers. *Scandinavian Journal of Information Systems*, *36*(2), 325–330.
- Aschenbrenner, L. (2024, May 29). Introduction SITUATIONAL AWARENESS: The Decade Ahead. https://situational-awareness.ai/
- aws. (2025). What is RLHF? Reinforcement Learning from Human Feedback Explained -AWS. Amazon Web Services, Inc. https://aws.amazon.com/what-is/reinforcementlearning-from-human-feedback/
- Belcic, I., & Stryker, C. (2024, September 18). What is GPT (generative pre-trained transformer)? | *IBM*. https://www.ibm.com/think/topics/gpt
- Bergmann, D. (2024, March 15). What is Fine-Tuning? | IBM. https://www.ibm.com/think/topics/fine-tuning
- Carroll, N., Holmström, J., Umeå University, Stahl, B. C., University of Nottingham, Fabian, N.
 E., & University of Groningen. (2024). Navigating the Utopia and Dystopia Perspectives of Artificial Intelligence. *Communications of the Association for Information Systems*, 55, 854–874. https://doi.org/10.17705/1CAIS.05533
- Chaudhari, S., Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K.,
 Deshpande, A., & Silva, B. C. da. (2024). RLHF Deciphered: A Critical Analysis of
 Reinforcement Learning from Human Feedback for LLMs (No. arXiv:2404.08555).
 arXiv. https://doi.org/10.48550/arXiv.2404.08555
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R.(2023). Economic potential of generative AI | McKinsey. *McKinsey Digital*.

https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economicpotential-of-generative-ai-the-next-productivity-frontier#

- Constantinides, P., Monteiro, E., & Mathiassen, L. (2024). Human-AI joint task performance: Learning from uncertainty in autonomous driving systems. *Information and Organization*, *34*(2). https://doi.org/10.1016/j.infoandorg.2024.100502
- Cui, T., Tan, B., & Shi, Y. (2024). Fostering humanistic algorithmic management: A process of enacting human-algorithm complementarity. *The Journal of Strategic Information Systems*, 33(2). https://doi.org/10.1016/j.jsis.2024.101838
- Curtis, A., & kidd, C. (2024). What Is Multimodal AI? A Complete Introduction. *Splunk*. https://www.splunk.com/en_us/blog/learn/multimodal-ai.html
- De Smet, A., Durth, S., Hancock, B., Mugayar-Baldocchi, M., & Reich, A. (2024). The Human Side of Generative AI: Creating a Path to Productivity. *Mckinsey & Company*. https://www.mckinsey.com/capabilities/people-and-organizational-performance/ourinsights/human-centered-ai-the-power-of-putting-people-first
- Demirci, O., Hannane, J., & Zhu, X. (2024). Research: How Gen AI Is Already Impacting the Labor Market. *Harvard Business Review*. https://hbr.org/2024/11/research-how-gen-ai-isalready-impacting-the-labor-market
- Dos Santos, J. F. P., & Williamson, P. J. (2024). Beyond connectivity: Artificial intelligence and the internationalisation of digital firms. *Information and Organization*, 34(4). https://doi.org/10.1016/j.infoandorg.2024.100538
- 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 (No. arXiv:2303.10130). arXiv. http://arxiv.org/abs/2303.10130

Feiner, L. (2024). Senator Laphonza Butler thinks supporting Big AI or human workers is a 'false choice.' *The Verge*. https://www.theverge.com/2024/10/3/24259479/senator-laphonza-butler-ai-california-workforce-harris

Financial Times. (2024a). Building the Generative AI advantage: Investing in skills & governance. *Fianncial Times - Partner Content*. https://www.ft.com/partnercontent/deloitte/building-the-generative-ai-advantage-

investing-in-skills-and-governance.html

Financial Times. (2024b). Transformative technology: What next for generative AI? Financial Times - Partner Content. https://www.ft.com/partnercontent/pwc/transformativetechnology-what-next-for-generative-ai.html

Frey, C. B., & Osborne, M. (n.d.). Generative AI and the Future of Work: A Reappraisal.

- Frey, C. B., & Osborne, M. (2024). Generative AI and the Future of Work: A Reappraisal. *Brown* Journal of World Affairs, 30.
- Fry, H. (Director). (2024, September 12). Super Intelligence [Broadcast]. https://www.bloomberg.com/news/videos/2024-09-12/super-intelligence-the-future-withhannah-fry-video
- Funk, J. (2024). What Is Normalization and When Should I Normalize? | mySidewalk Knowledge Center. https://help.mysidewalk.com/en/articles/3126175-what-isnormalization-and-when-should-i-normalize
- Goasduff, L. (2024). Gartner Says Generative AI will Require 80% of Engineering Workforce to Upskill Through 2027. *Gertner Newsroom*. https://www.gartner.com/en/newsroom/pressreleases/2024-10-03-gartner-says-generative-ai-will-require-80-percent-of-engineeringworkforce-to-upskill-through-2027

- Google. (2025). Overview of GAN Structure | Machine Learning. *Google for Developers*. https://developers.google.com/machine-learning/gan/gan_structure
- Gordon, R. (2023, March 3). Large language models are biased. Can logic help save them? MIT News | Massachusetts Institute of Technology. https://news.mit.edu/2023/large-languagemodels-are-biased-can-logic-help-save-them-0303
- Gruber, T. (2017). Tom Gruber: How AI can enhance our memory, work and social lives | *TED Talk*.

https://www.ted.com/talks/tom_gruber_how_ai_can_enhance_our_memory_work_and_s ocial_lives

- Hadidi, R., & Klein, B. (2024). The Future of Work, Physical Location of Workers,
 Technological Issues and Implications. *Journal of the Midwest Association for Information Systems*, 2024(1).
- Hall, B. (2017, June 1). Incorporating Non-Academic Sources in Research Articles. Sengi. http://15.157.18.82/research-tips/57-incorporating-non-academic-sources-in-researcharticles/
- Hamirani, Q. (2024). Here's How Generative AI Will Redefine The Workplace. *Forbes*. https://www.forbes.com/sites/qhamirani/2024/02/15/heres-how-generative-ai-will-redefine-the-workplace/
- Hazan, E., Madgavkar, A., Chui, M., Smit, S., Maor, D., Dandona, G. S., & HuyghuesDespointes, R. (2024). A new future of work: The race to deploy AI and raise skills in
 Europe and beyond. *Mckinsey Global Institute*. https://www.mckinsey.com/mgi/ourresearch/a-new-future-of-work-the-race-to-deploy-ai-and-raise-skills-in-europe-andbeyond

- Holmström, J. (2024). A Layered Organizing Logic for Generative AI Evolution. Insights from digital infrastructure theory. *36*(1), 11.
- Hoover, A. (2023). Use of AI Is Seeping Into Academic Journals—And It's Proving Difficult to Detect. *Wired*. https://www.wired.com/story/use-of-ai-is-seeping-into-academic-journals-and-its-proving-difficult-to-detect/
- Huy, L. V., Nguyen, H. T. T., Vo-Thanh, T., Thinh, N. H. T., & Dung, T. T. T. (2024). Generative AI, Why, How, and Outcomes: A User Adoption Study. *AIS Transactions on Human-Computer Interaction*, 16(1), 1–27. https://doi.org/10.17705/1thci.00198
- IBM. (2023, October 12). Understanding the Different Types of Types of Artificial Intelligence. https://www.ibm.com/think/topics/artificial-intelligence-types
- Illinois.edu. (2024). Traditional AI vs. Generative AI: What's the Difference? *College of Education*. https://education.illinois.edu/about/news-events/news/2024/11/11/what-is-generative-ai-vs-ai
- Jarrahi, M. H. (2018). Artificial Intelligence and the Future of Work: Human-AI Symbiosis in Organizational Decision Making. *Business Horizons*, 61. https://doi.org/10.1016/j.bushor.2018.03.007
- Keen, E. (2024). Gartner Predicts 40% of Generative AI Solutions Will Be Multimodal By 2027. Gartner Newsroom. https://www.gartner.com/en/newsroom/press-releases/2024-09-09gartner-predicts-40-percent-of-generative-ai-solutions-will-be-multimodal-by-2027
- Kinder, M., de Souza Briggs, X., Muro, M., & Liu, S. (2024, October 10). Generative AI, the American worker, and the future of work. *Brookings*. https://www.brookings.edu/articles/generative-ai-the-american-worker-and-the-future-ofwork/

- Klingler, N. (2024, February 13). The 3 Types of Artificial Intelligence: ANI, AGI, and ASI. *Viso.Ai.* https://viso.ai/deep-learning/artificial-intelligence-types/
- Knight, W. (2024). No One Actually Knows How AI Will Affect Jobs. *Wired*. https://www.wired.com/story/ai-impact-on-work-mary-daly-interview/
- Kuusi, O., & Heinonen, S. (2022). Scenarios From Artificial Narrow Intelligence to Artificial General Intelligence—Reviewing the Results of the International Work/Technology 2050 Study. *World Futures Review*, 14(1), 65–79. https://doi.org/10.1177/19467567221101637
- Lanamäki, A., Väyrynen, K., Hietala, H., Parmiggiani, E., & Vassilakopoulou, P. (2024). Not Inevitable: Navigating Labor Displacement and Reinstatement in the Pursuit of AI for Social Good. *Communications of the Association for Information Systems*, 55, 831–845. https://doi.org/10.17705/1CAIS.05531
- Lang, G., Triantoro, T., & Sharp, Ja. (2024). Large Language Models as AI-Powered Educational Assistants: Comparing GPT-4 and Gemini for Writing Teaching Cases. *Journal of Information Systems Education*, 35(3), 390–407. https://doi.org/10.62273/YCIJ6454
- Larson, L. (2024). AI in the workplace, and a defense of busywork. *The Verge*. https://www.theverge.com/24066270/ai-automation-work-labor-busywork
- lascaze, E., Corduneanu, R., Kreit, B., Cantrell, S., Kulkarni, A., & Rifkin, D. (12/24). Entry level jobs reskilling for AI. *Deloitte Insights*. https://www2.deloitte.com/us/en/insights/topics/talent/ai-in-the-workplace.html
- Li, C., Wang, H., Jiang, S., & Gu, B. (2024). The Effect of AI-Enabled Credit Scoring on Financial Inclusion: Evidence from an Underserved Population of over One Million. *MIS Quarterly*, 48(4), 1803–1834. https://doi.org/10.25300/MISQ/2024/18340

Li, Y., Datta, S., Rastegar-Mojarad, M., Lee, K., Paek, H., Glasgow, J., Liston, C., He, L., Wang, X., & Xu, Y. (2025). Enhancing systematic literature reviews with generative artificial intelligence: Development, applications, and performance evaluation. *Journal of the American Medical Informatics Association*, ocaf030.

https://doi.org/10.1093/jamia/ocaf030

- Livari, N., Sharma, S., Sharma, S., Hartikainen, H., & Ventä-Olkkonen, L. (2024). From Understanding AI Literacy to Fostering Children's Growth. Towards powerful change makers regarding intelligent, data-driven technologies. *Scandinavain Journal of Information Systems*, 36(2).
- Lovich, D., Sargeant, R., & Smith, J. (2024, June 12). How Gen AI Can Make Work More Fulfilling. *Harvard Business Review*. https://hbr.org/2024/06/how-gen-ai-can-makework-more-fulfilling
- Lunden, I. (2025, January 15). LinkedIn adds free AI tools for job hunters and recruiters. *TechCrunch*. https://techcrunch.com/2025/01/15/linkedin-ads-free-ai-tools-for-job-hunters-and-recruiters/
- Ma, H. V., Huang, W. W., & Dennis, A. R. (2024). Unintended Consequences of Disclosing Recommendations by Artificial Intelligence versus Humans on True and Fake News Believability and Engagement. *Journal of Management Information Systems*, 41(3), 616– 644. https://doi.org/10.1080/07421222.2024.2376381
- Mahedi Hasan, H. M., Sanyal, F., Chaki, D., & Ali, Md. H. (2017). An empirical study of important keyword extraction techniques from documents. 2017 1st International Conference on Intelligent Systems and Information Management (ICISIM), 91–94. https://doi.org/10.1109/ICISIM.2017.8122154

Marcus, G. (2024). Generative AI Still Needs to Prove Its Usefulness. *Wired*. https://www.wired.com/story/generative-ai-will-need-to-prove-its-usefulness/

- Marr, B. (2023). A Short History Of ChatGPT: How We Got To Where We Are Today. *Forbes*. https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-howwe-got-to-where-we-are-today/
- Marr, B. (2024, January 4). Generative AI And The Future Of Jobs. https://www.forbes.com/sites/bernardmarr/2024/01/04/generative-ai-and-the-future-ofjobs/
- Maurya, H., Agrahari, A., & Kumar, A. (2024). Generative AI and Future of Work: Enhancing Customer's Experience of a Ridehailing Application. *Communication of the Association for Information Systems*, 55, 459–473. https://doi.org/10.17705/1CAIS.05518
- Mayer, H., Yee, L., Chui, M., & Roberts, R. (2025). Superagency In the Workplace: Empowering people to unlock AI's full potential. *Mckinsey & Company*.
 https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/gen-ai-a-cognitive-industrial-revolution
- McDermott, K. (2024). Knowledge workers are leaning on generative AI as their workloads mount. *VentureBeat*. https://venturebeat.com/programming-development/knowledgeworkers-are-leaning-on-generative-ai-as-their-workloads-mount/
- McDonough, M. (2024). Precision and recall | F-score, Formula, & Facts | Britannica. *Brittanica*. https://www.britannica.com/science/precision-and-recall
- McKendrick, J. (2024). How Generative AI Rattles the Workplace. *Forbes*. https://www.forbes.com/sites/joemckendrick/2024/07/12/how-generative-ai-rattles-theworkplace/

McKinsey. (2024). What is prompt engineering? | McKinsey.

https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-promptengineering

- McLean, S., Read, G. J. M., Thompson, J., Baber, C., Stanton, N. A., & Salmon, P. M. (2023).
 The risks associated with Artificial General Intelligence: A systematic review. *Journal of Experimental & Theoretical Artificial Intelligence*, *35*(5), 649–663.
 https://doi.org/10.1080/0952813X.2021.1964003
- Miller, R. (2024, June 19). In spite of hype, many companies are moving cautiously when it comes to generative AI. *TechCrunch*. https://techcrunch.com/2024/06/19/in-spite-ofhype-many-companies-are-moving-cautiously-when-it-comes-to-generative-ai/
- MIT Insights. (2024). Moving generative AI into production. *MIT Technology Review Insights*. https://www.technologyreview.com/wp-

content/uploads/2024/11/MITTR_Redis_final_26nov24

Mittal, N., Perricos, C., Sterrett, L., & Dutt, D. (2024). Generative AI and the future of work: The Potential? Boundless. *Deloitte AI Institute*.

https://www2.deloitte.com/us/en/pages/consulting/articles/gen-ai-use-cases.html

- Mozelius, P., & Humble, N. (2024). On the Use of Generative AI for Literature Reviews: An Exploration of Tools and Techniques. *European Conference on Research Methodology* for Business and Management Studies, 23(1), 161–168. https://doi.org/10.34190/ecrm.23.1.2528
- Muro, M., Maxim, R., & Whiton, J. (2019). Automation and Artificial Intelligence: How machines are affecting people and places. *Brookings*.

https://www.brookings.edu/articles/automation-and-artificial-intelligence-how-machinesaffect-people-and-places/

- Nicoletti, L., & Bass, D. (2023). Generative AI Takes Stereotypes and Bias From Bad to Worse. https://www.bloomberg.com/graphics/2023-generative-ai-bias/
- NNLM. (2025). Generative Artificial Intelligence. *NNLM*. https://www.nnlm.gov/guides/data-thesaurus/generative-artificial-intelligence
- Nyman, S., Benfeldt, O., Nørbjerg, J., & Jensen, T. B. (2024). Towards a Research Agenda on Data-Driven Labour Organizing. *Scandinavian Journal of Information Systems*, 36(2), 317–322.
- O'Connor, R. (2022). Introduction to Diffusion Models for Machine Learning. *AssemblyAI*. https://assemblyai.com/blog/diffusion-models-for-machine-learning-introduction
- OpenAI. (2024a). Accelerating the development of life-saving treatments. *OpenAI*. https://openai.com/index/moderna/

OpenAI. (2024b). Creating a GPT | OpenAI Help Center.

https://help.openai.com/en/articles/8554397-creating-a-gpt

- OpenAI. (2024c). Delivering contextual job matching for millions with OpenAI. *OpenAI*. https://openai.com/index/indeed/
- OpenAI. (2024d). Partnering with Axios expands OpenAI's work with the news industry. *OpenAI*. https://openai.com/index/partnering-with-axios-expands-openai-work-with-thenews-industry/
- Pasch, H. (2024). Generative AI isn't coming for you—Your reluctance to adopt it is. *VentureBeat*. https://venturebeat.com/ai/generative-ai-isnt-coming-for-you-yourreluctance-to-adopt-it-is/

- Plumb, T. (2024). Tech giants: AI will transform 92% of ICT jobs; we must upskill now. *VentureBeat.* https://venturebeat.com/ai/tech-giants-ai-will-transform-92-of-ict-jobs-wemust-upskill-now/
- Poitevin, H. (2024). AI in the Workplace: How Generative AI Is Transforming Workforce Roles and HR Strategies. *Gartner*. https://www.gartner.com/en/videos/ai-in-the-workplace
- Potter, N. (2024, February 27). AI May Create Far More Jobs Than It Kills. *IEEE Spectrum*. https://spectrum.ieee.org/ai-fears-jobs-technology
- Pretz, K. (2025). AI leads the way in the 2025 IEEE Tech Impact Study. *IEEE Spectrum*. https://spectrum.ieee.org/2025-ieee-tech-impact-study
- PwC. (2024a). Enter the future of work with AI: A new era of "agentic" work has begun. PwC Insights. https://www.pwc.com/us/en/services/ai/ai-and-the-future-of-work.html
- PwC. (2024b). How PwC is using GenAI to transform business value. PwC Insights. https://www.pwc.com/us/en/tech-effect/ai-analytics/generative-ai-impact-onbusiness.html
- PwC. (2024c). Leveraging Generative AI for Job Augmentation and Workforce Productivity: Scenarios, Case Studies and a Framework for Action. *PwC Insights*. https://www.pwc.com/gx/en/issues/artificial-intelligence/wef-leveraging-generative-aifor-job-augmentation-and-workforce-productivity-2024.pdf
- Queiroz, M., Anand, A., & Baird, A. (2024). Manager Appraisal of Artificial Intelligence Investments. *Journal of Management Information Systems*, 41(3), 682–707. https://doi.org/10.1080/07421222.2024.2376383
- Rahman, R. (2024, June 19). The training compute of notable AI models is doubling roughly every five months. *Epoch AI*. https://epoch.ai/data-insights/compute-trend-post-2010

- Rosani, G., & Farri, E. (2024, September 20). How the Next Generation of Managers Is Using Gen AI. *Harvard Business Review*. https://hbr.org/2024/09/how-the-next-generation-ofmanagers-is-using-gen-ai
- Roslansky, R. (2024). The AI-Fueled Future of Work Needs Humans More Than Ever. *Wired*. https://www.wired.com/story/ai-fueled-future-of-work-needs-humans-more-than-ever/

Rotman, D. (2024a, January 27). People are worried that AI will take everyone's jobs. We've been here before. *MIT Technology Review*.
https://www.technologyreview.com/2024/01/27/1087041/technological-unemployment-elon-musk-jobs-ai/

- Rotman, D. (2024b, August 20). How to fine-tune AI for prosperity. *MIT Technology Review*. https://www.technologyreview.com/2024/08/20/1096733/how-to-fine-tune-ai-forprosperity/
- Saffarizadeh, K., Keil, M., & Maruping, L. (2024). Relationship Between Trust in the AI Creator and Trust in AI Systems: The Crucial Role of AI Alignment and Steerability. *Journal of Management Information Systems*, 41(3), 645–681.

https://doi.org/10.1080/07421222.2024.2376382

- Schreiner, M. (2023, July 11). GPT-4 architecture, datasets, costs and more leaked. *THE DECODER*. https://the-decoder.com/gpt-4-architecture-datasets-costs-and-more-leaked/
- Shact, L., Kreit, B., Vert, G., Holdowsky, J., & Buckley, N. (2024). Four futures of generative AI in the enterprise: Scenario planning for strategic resilience and adaptability. *Deloitte Insights*. https://www2.deloitte.com/us/en/insights/topics/digitaltransformation/generative-ai-and-the-future-enterprise.html

- Shahid, A., & Mishra, S. (2024). A Framework for a Master's in Applied Artificial Intelligence Program in Computer and Information Systems Discipline. *Journal of Information Systems Education*, 35(4), 495–511. https://doi.org/10.62273/EQZE3625
- Singla, A., Sukharevsky, A., Yee, L., Chui, M., & Hall, B. (2025, March 12). The State of AI: Global survey | McKinsey. *QuantumBlack*.

https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. Journal of Business Research, 104, 333–339. https://doi.org/10.1016/j.jbusres.2019.07.039

Stohr, A., Ollig, P., Keller, R., & Rieger, A. (2024). Generative mechanisms of AI implementation: A critical realist perspective on predictive maintenance. *Information and Organization*, 34(2). https://doi.org/10.1016/j.infoandorg.2024.100503

Stryker, C. (2024, July 15). What is Multimodal AI? | IBM. https://www.ibm.com/think/topics/multimodal-ai

- Suleyman, M. (2023). The Coming Wave: Technology, power, and the twenty-first century's greatest dilemma. *Crown*.
- Sundberg, L., & Holmström, J. (2024a). Fusing domain knowledge with machine learning: A public sector perspective. *The Journal of Strategic Information Systems*, 33(3). https://doi.org/10.1016/j.jsis.2024.101848
- Sundberg, L., & Holmström, J. (2024b). Teaching Tip Using No-Code AI to Teach Machine Learning in Higher Education. *Journal of Information Systems Education*, 56–66. https://doi.org/10.62273/CYPL2902

Taloka.ai. (2023). History of generative AI. https://toloka.ai/blog/history-of-generative-ai/

 Tao, J., Zhou, L., & Fang, X. (2024). Generative AI for Intelligence Augmentation:
 Categorization and Evaluation Frameworks for Large Language Model Adaptation. *AIS Transactions on Human-Computer Interaction*, *16*(3), 364–387.
 https://doi.org/10.17705/1thci.00210

- Teli, J. S., Rai, A., & Lin, Y.-K. (2024). Abnormal Returns to Artificial Intelligence Patent Infringement Litigations. *Journal of Management Information Systems*, 41(2), 422–452. https://doi.org/10.1080/07421222.2024.2340826
- The Week. (2024, July 31). Future of generative AI: Utopia, dystopia or up to us? *The Week*. https://theweek.com/tech/future-of-generative-ai-utopia-dystopia-or-up-to-us
- Thornhill, J. (2024). Superfluous people vs AI: What the jobs revolution might look like. *Financial Times*. https://www.ft.com/content/f9b964bd-0edc-4dfd-ad30-88b011cd4fde
- Tutun, S., Topuz, K., Tosyali, A., Bhattacherjee, A., & Li, G. (2024). Explainable Artificial Intelligence for Mental Disorder Screening: A Computational Design Science Approach. *Journal of Management Information Systems*, 41(4), 958–981. https://doi.org/10.1080/07421222.2024.2415771
- Tzouganatou, A. (2018). Can Heritage Bots Thrive? Toward Future Engagement in Cultural Heritage. Advances in Archaeological Practice, 6(4), 377–383. https://doi.org/10.1017/aap.2018.32
- University College London. (2024a). AIs are irrational, but not in the same way that humans are. *ScienceDaily*. https://www.sciencedaily.com/releases/2024/06/240605162628.htm
- University College London. (2024b). Bias in AI amplifies our own biases, researchers show. *ScienceDaily*. https://www.sciencedaily.com/releases/2024/12/241218132137.htm

Urwin, M. (2024). Precision and Recall in Classification Models. *Built In*. https://builtin.com/data-science/precision-and-recall

- Wang, Q., Liu, X., & Huang, K.-W. (2024). Investigating employees' occupational risks and benefits resulting from artificial intelligence: An empirical analysis. *Information & Management*, 61(8), 104036. https://doi.org/10.1016/j.im.2024.104036
- Warren, T. (2024). Microsoft's new Copilot AI agents act like virtual employees to automate tasks. *The Verge*. https://www.theverge.com/2024/5/21/24158030/microsoft-copilot-ai-automation-agents
- Wiggers, K. (2025). Generative AI funding reached new heights in 2024. *TechCrunch*. https://techcrunch.com/2025/01/03/generative-ai-funding-reached-new-heights-in-2024/
- Woodall, T. (2024, May 23). AI poised to usher in new level of concierge services to the public. *ScienceDaily*. https://www.sciencedaily.com/releases/2024/05/240523153643.htm
- Zorpette, G. (2024, June 6). How Large Language Models Are Changing My Job. *IEEE* Spectrum. https://spectrum.ieee.org/large-language-models-2668430044
- Zysman, J., & Nitzberg, M. (2024). Generative AI and the Future of Work: Augmentation or Automation? *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4811728

Appendix 1: Cyclical Training Process Example with Dialog

A new conversation is created with the custom GPT, training article is uploaded. The GPT is prompted to follow *Instructional Prompt 2* (see Methodology 3.2.B) and its knowledge documents to analyze the article. The initial output was documented for future reference, and the custom GPT receives direct reinforcement training. To begin this process, *Training Prompt 1* is entered to the chat as follows:

Training Prompt 1:

"I am now going to provide you feedback on your analysis. I want you remember this feedback I give you going forward into our next analysis. At any time, you are confused, I want you to say "Need help". "

The custom GPT was then prompted to examine the response for every topic one at a time by re-generating the exact quotes from the paper. After the exact quotes for the first topic were generated, the GPT was asked to provide what it believes constitutes the presence of that topic. This process served to provide knowledge about the GPT's reasoning as well as confirmation of understanding. Feedback was then provided for each topic, providing feedback on accuracy and definitions. Below is an excerpt from the training dialog with AI generated text shown italicized:

Custom GPT:

Certainly! Here are the quotes I provided for the Augmentation topic: 1. "The ability to adopt generative AI into their workflow will make them more efficient, productive and valuable."

2. "But as well as enhancing our technical skills, delegating routine elements of work – scheduling, drafting reports – to machines means we have more time to focus on the human aspect of our work."

Trainer:

I like your analysis here. However, I am concerned that you did not evaluate the whole document for instances of augmentation sentiment. I believe you found two instances and did not continue to search. I believe this to be true because you did not find the line "there's a much higher chance you'll use AI to augment your work.". Going forward, I would like you to identify all the instances of the sentiment of each topic.

Custom GPT:

Thank you for the feedback and for pointing out the missed instance. You're absolutely right—my analysis should capture all instances of a topic, and the quote "there's a much higher chance you'll use AI to augment your work" indeed represents augmentation sentiment. I acknowledge that I stopped prematurely, and I will ensure to fully evaluate the document for every occurrence of a topic moving forward.

I'll take extra care to thoroughly review each document, identify every relevant instance, and include all appropriate quotations. This will improve the comprehensiveness and accuracy of my analyses. Thank you for emphasizing this!

This process was repeated for each topic, with the trainer providing feedback on the original output. The goal was to provide the GPT with enhanced reasoning and comprehension abilities for future analyses. After each topic was reinforced, the paper was again uploaded to the conversation and asked to follow its instructional prompt using its knowledge documents. This process was repeated for until the custom GPT generated an output identical to the human analysis. Once the generated output was as accurate as the human analysis, the custom GPT was prompted to summarize the conversation with emphasis on training feedback using *Training Summary Prompt 1* shown below. The GPT generated response is shown below the prompt in italicized text.

Training Summary Prompt 1:

"I want to summarize the training I have given you. Could you put into words all the feedback I have given you and what you have learned"

Custom GPT:

Absolutely! Here's a summary of the key feedback you've provided during our training and the adjustments I've made to refine my analysis process. This can serve as a reinforcement learning log to encapsulate our progress: Reinforcement Training Summary

General Approach

1. Thorough Analysis:

•Always examine the entire document for instances of the specified topics. Avoid prematurely concluding a topic's analysis and ensure all relevant sentiments are captured.

2. Contextual Sensitivity:

•Interpret quotes with attention to nuanced meanings. Ensure that quotes chosen for a topic directly reflect its sentiment and do not overlap without justification.

Topic-Specific Insights

1.Automation:

•Automation sentiment should focus on task replacement or minimal human involvement due to AI. Avoid conflating this with augmentation unless explicit indications of replacement are present.

•Feedback Applied: Exclude quotes that suggest transformation but do not imply elimination of jobs.

2. *Augmentation*:

•Augmentation refers to enhancing productivity or freeing humans for strategic or creative work. Be diligent in identifying keywords like "enhanced" or "productive" and ensure all mentions are included.

•Feedback Applied: Capture all instances of augmentation, such as mentions of workflows, creativity, or human-centric improvements. 3.Hallucination:

Hallucination involves the AI system producing false or fabricated outputs. Focus on explicit mentions of inaccuracy, misinformation, or risks of AI creating incorrect content.
Feedback Applied: Confirmed hallucination is absent when not directly mentioned.
4.Bias:

•Bias includes both systemic biases in AI outputs and broader societal inequalities exacerbated by AI. Distinguish societal effects from technical system biases, but now include societal biases per updated guidance.

•*Feedback Applied: Reassessed quotes to ensure both dimensions of bias are captured appropriately.*

5.Regulatory:

Regulatory concerns address compliance, governance, and legal standards related to AI systems. These include data privacy laws and ethical accountability frameworks.
Feedback Applied: Confirmed regulatory absence when no legal or governance frameworks were mentioned.

6.*Ethics*:

•Ethical discussions often overlap with bias but focus on responsibility, fairness, and human well-being. Ensure examples reflect moral implications, equitable opportunities, and preservation of human-centric work.

•*Feedback Applied: Broadened ethical analysis to include fairness in AI-driven societal changes.*

7.Dystopian Futures (Doomsday):

•Look for explicit mentions of existential risks or speculative worst-case scenarios involving AI. Avoid including vague concerns or unrelated challenges.

•Feedback Applied: Confirmed absence of doomsday mentions when not present.

Key Learning Points

1.Relevance and Precision:

•Only include quotes that clearly align with the sentiment of the topic. Avoid overgeneralizing or misinterpreting ambiguous language.

2. Incorporating Feedback Dynamically:

•Adjust topic definitions and interpretations based on guidance, such as expanding bias to include societal inequalities.

3. Holistic Document Review:

•Ensure that every instance of a topic is captured without skipping or prematurely concluding the analysis.

4. Iterative Refinement:

•*Reassess earlier decisions to integrate clarified definitions or updated topic scopes, ensuring consistency and accuracy.*

Outcome

Through this training, I have refined my analytical rigor, improved contextual sensitivity, and gained a deeper understanding of nuanced distinctions between topics. This iterative process has strengthened my ability to adapt dynamically to feedback and produce more comprehensive evaluations.

This output summarizes the reinforcement training provided to the model. This output was saved to a word document and uploaded to the custom GPT's knowledge configuration to be used as reference for future analyses. This process was repeated for all training papers as part of the cyclical training process.

Appendix 2: Statement Regarding Generative AI Use

All use of Generative AI (GAI) in this thesis complies with the Honors Tutorial College Thesis Guidelines. Formal approval for the use of GAI in the described manner was obtained from both the HTC Dean and my Director of Studies on November 1, 2024. Below is Subsection 2 and 3 of the HTC generative AI Guidelines:

2. Using AI to generate a literature review without citing the software's work in generating sources.

This is unacceptable use because students must base their work on substantial research using the acknowledged standards of their field; thus, using AI in this manner would negate their fulfillment of the assignment.

Generative AI was used to generate a literature review as a core methodological contribution of my thesis. This was extensively cited, and my work was based on substantial research using the standards of the AIS field.

3. Using AI to generate arguments or interpret sources (primary or secondary).

This is unacceptable use because students must offer an innovative interpretation of important issues or ideas in their field. If AI generates this work, then the student does not fulfill this criterion.

At no point was Generative AI used to construct arguments or interpret sources beyond the scope explicitly defined in the methodology.