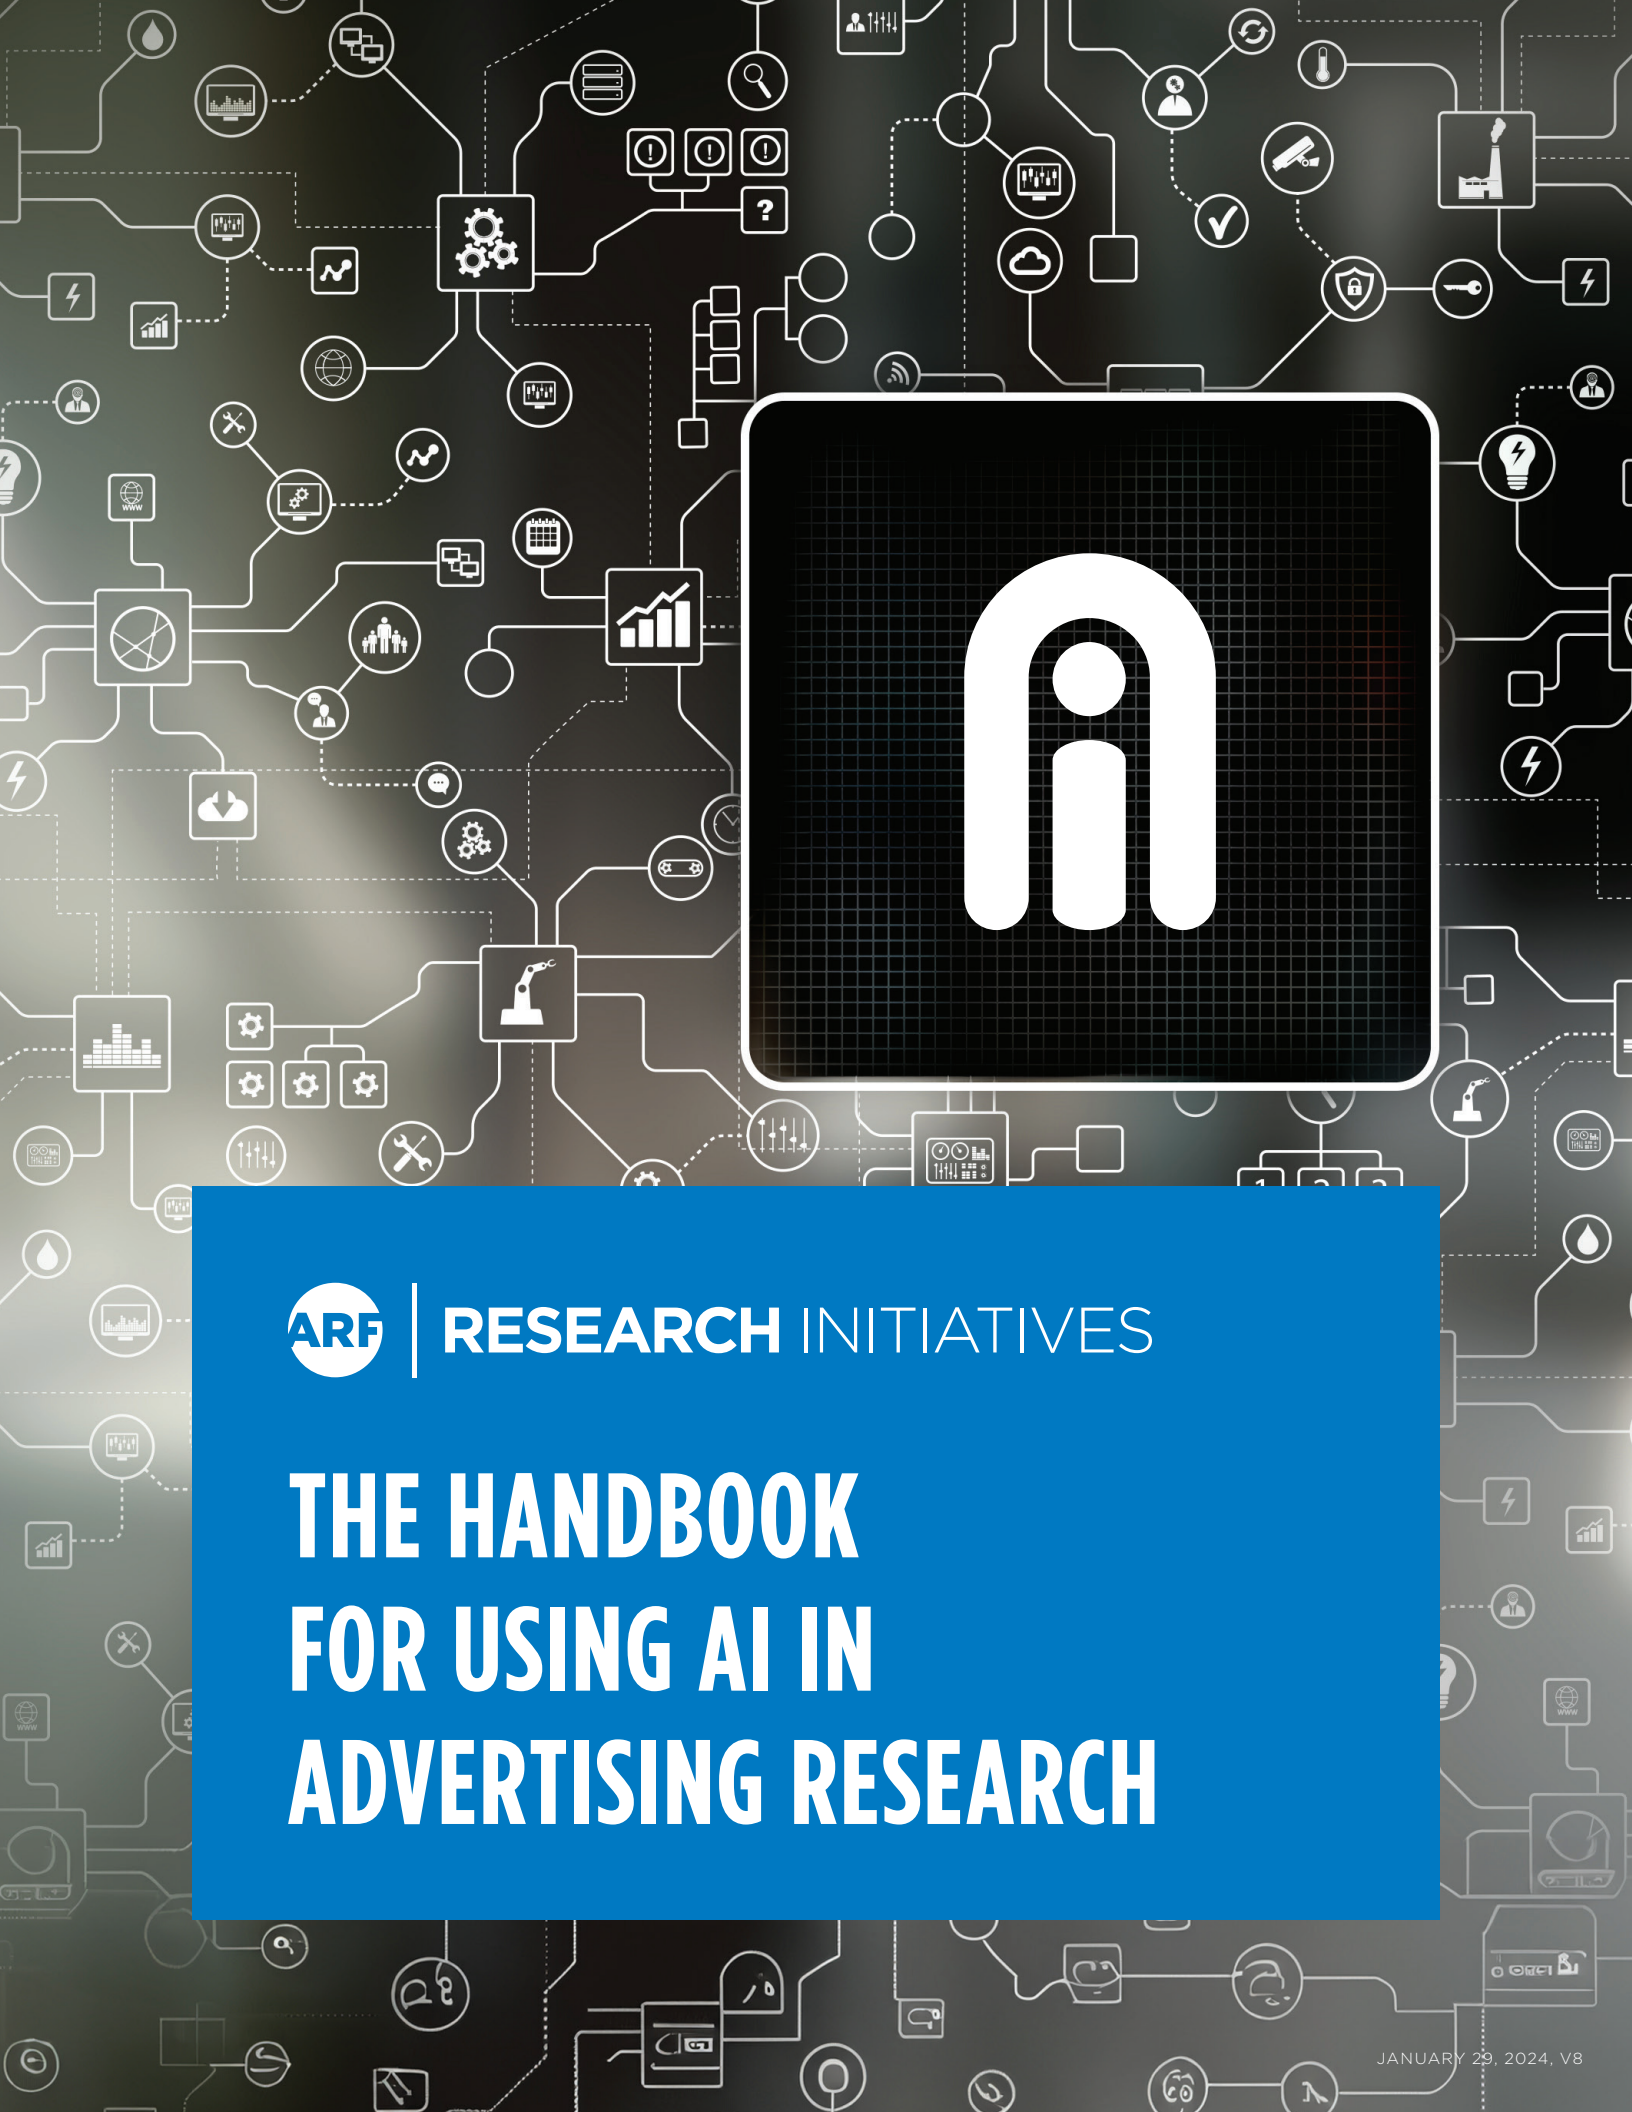

ARF | RESEARCH INITIATIVES

THE HANDBOOK FOR USING AI IN ADVERTISING RESEARCH

JANUARY 29, 2024, V8

CO-AUTHORS

Kelly Castellon
GROUP M

Jackie Day
GROUP M

Nicole Mitchell
DYNATA

Kiko Restrepo
WARNER BROTHERS DISCOVERY

Laura Vanison
VEVO

Agustina Perez Blua
(YOUNG PROS OFFICER)
GOOGLE

ARF CULTURAL EFFECTIVENESS COUNCIL CURRENT MEMBERS

Janelle James
(CO-CHAIR)
IPSOS

Barbara Leflein
(OUTGOING CO-CHAIR)
LEFLEIN ASSOCIATES

Dana Sparber
(CO-CHAIR)
NBCUNIVERSAL

Alexa Blechman
A+E NETWORKS

Kelly Castellon
GROUPM

Jackie Day
GROUPM

Donna Dei-Baning
DENTSU X, LLC

Tony D'Andrea
GENERAL MILLS

Daniela Iglesias
MAGID

Jonathan Lopes
DIGITAS

Joel Marquez
CHARTER COMMUNICATIONS OPERATING, LLC

Tristan Marra
GLAAD

Nicole Mitchell
DYNATA

Karen Ramspacher
MRI SIMMONS

Kiko Restrepo
WARNER MEDIA

Chloe Stromberg
LINKEDIN

Laura Vanison
VEVO

Tarya Weedon
HOROWITZ RESEARCH

Anna Wilgan
KANTAR

Anne Kaplan
(YOUNG PROS OFFICER)
PARAMOUNT

Agustina Perez Blua
(YOUNG PROS OFFICER)
GOOGLE

Danielle Zito
(YOUNG PROS OFFICER)
IPSOS

CONTENTS

INTRODUCTION

The integration of Generative AI (gen AI) is rapidly growing in businesses (Chui et al., 2023). With many organizations already incorporating gen AI in various functions (see Ipsos Knowledge Panel Survey, 2023; Mittani et al., 2022) AI has shifted from a technical subject to a strategic business concern, with a notable percentage of C-suite executives and boards now engaging with gen AI tools. Additionally, a considerable number of companies plan to increase their AI investment due to gen AI advancements. According to WARC's Marketer's Toolkit (2023) 58% of respondents identify themselves as "cautiously progressive" towards gen AI, actively engaging in trials with large language model chatbots for strategic insights and exploring text-to-image AI applications to enhance creativity. Alongside this potential and growth, there are abundant challenges in managing gen AI-related risks, particularly around inaccuracy, with many organizations still developing strategies for mitigation.

In the ever-evolving landscape of advertising research, artificial intelligence (AI) could stand as a transformative force, reshaping how data is analyzed, insights are gleaned, and strategies are formed. "The ARF Handbook for Using AI in Advertising Research" is a comprehensive guide designed to navigate the complexities and potentials of AI in this dynamic field. This handbook aims to demystify AI and its myriad applications in advertising research, offering a clear path for professionals to integrate these technologies into their work effectively.

From the foundational concepts in *Chapter 1:The Evolution of AI* and *Chapter 2: Basics of AI and Machine Learning* to the comprehensive literature review focusing on practical applications highlighted in *Chapter 3: Integrating AI into Marketing and Advertising Research*, this handbook provides a thorough grounding in AI's principles and practices. The heart of the handbook lies in *Chapter 4: ARF Case Studies – AI Implementations in Advertising Research*, presenting real-world scenarios ranging from survey design to synthetic research, responding to the existing lacuna in AI utilization in advertising research and demonstrating AI's potentially profound impact on the industry. These case studies compare the relative performance of three of the most widely used LLMs: ChatGPT 4, Bard and Claude AI. They are presented here as concise versions with links provided for readers who wish to dive more deeply into the intricacies of any of the case studies. Chapter 5 delves into issues of *Ethics and Transparency in AI-Driven Advertising Research*, underscoring the need for practitioners to uphold the highest standards of privacy, fairness, and clarity. Finally, looking ahead, *Chapter 6: The Future of AI in Advertising Research* explores emerging trends, challenges, and opportunities, equipping readers to stay ahead in a rapidly changing domain.

This handbook is not just a repository of knowledge but a call to action – to embrace the AI revolution in advertising research and continue learning and adapting in this exciting journey. We thus open the possibility for ARF members to be active contributors by creating this handbook in the form of a WIKI. Your experiences, insights, and knowledge are invaluable in enriching this resource. **Please visit the AI page on the ARF website to add your perspective and be a part of this research.** Your input makes a difference!

CHAPTER 1: THE EVOLUTION OF AI

1. THEORETICAL FOUNDATIONS:

a. Pre-1950s: Philosophers and mathematicians, like Alan Turing, ponder the idea of machines that think. Turing's 1950 seminal paper, "Computing Machinery and Intelligence," introduced the 'imitation game,' or *Turing Test* to determine machine intelligence. Offering a criterion for thinking, Turing proposed to consider the question "Can machines think?" using this test principle: If x plays Turing's imitation game satisfactorily, then x thinks. He states the following:

> I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9, to make them play the imitation game so well that an average interrogator will not have more than 70 per cent, chance of making the right identification after five minutes of questioning. The original question, 'Can machines think!' I believe to be too meaningless to deserve discussion. Nevertheless I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted. (p. 442).

Figure 1: Milestones in computational pathology. Bera et al., (2019)

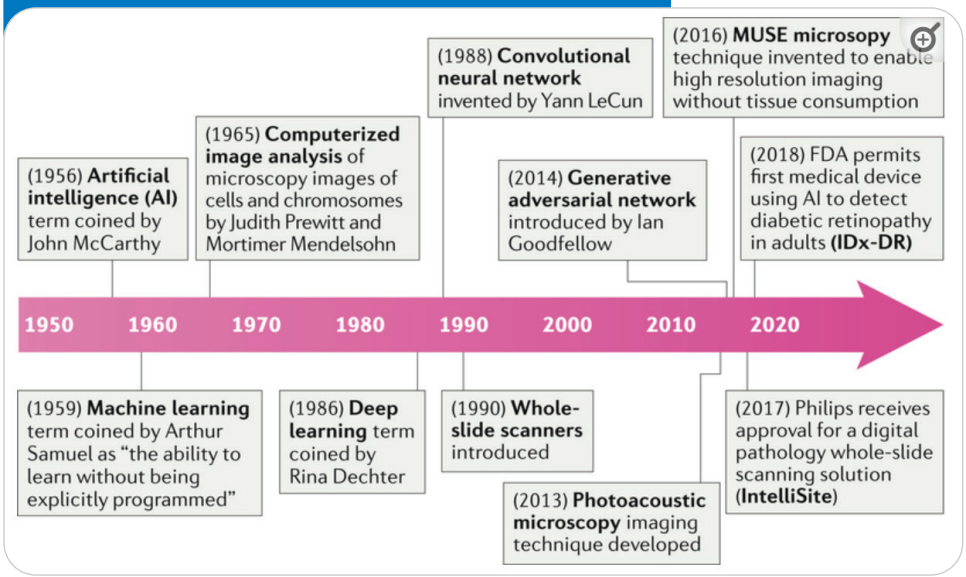

2. BIRTH OF AI:

a. 1956: The Dartmouth Workshop takes place in Hanover, New Hampshire. The term "Artificial Intelligence" is coined by John McCarty, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, and the foundational goals of the field are established (McCarthy et al., 2006). AI is referred as the branch of computer science where "machine-based approaches are used to make a prediction – emulating what an intelligent human might do in the same situation" (p. 2).

b. Late 1950s to 1960s: Early AI research leads to the creation of programs that mimic human abilities, such as problem-solving and playing chess. In the 1950s, AI research delved into areas such as problem-solving and symbolic approaches. The subsequent programs that emerged post the Dartmouth Workshop left most observers awe inspired. Computers were suddenly capable of solving complex algebraic word problems, proving theorems in geometry, and even acquiring the ability to communicate in English. This era of research was marked by fervent optimism, with many experts confidently predicting the creation of a fully intelligent machine within two decades (Synaptiq, 2020). As the 1960s rolled in, the United States Department of Defense recognized the potential of this burgeoning field and initiated efforts to train computers in emulating rudimentary human reasoning. For instance, the Defense Advanced Research Projects Agency (DARPA) successfully executed street mapping projects during the 1970s as part of these endeavors (Naimark, 2006; Lippman, 1980). Goldstein and Papert (1977) characterized the exhibitions of the Dendral program (Lindsay et al., 1980) (the first expert system that automated decision-making process and problem-solving behavior) in the mid-1960s as a "paradigm shift" within the field of AI, marking a transition towards knowledge-based systems (Buchanan, 2005).

3. AI WINTER AND FUNDING CHALLENGES:

a. Early 1970s and Late 1980s: Initial optimism about AI led to high expectations, but technological limitations led to disappointment and reduced funding, known as the "AI winter." During the 1970s, AI faced criticism and financial setbacks as researchers struggled to tackle increasingly complex problems. A major obstacle was the limited computational power available at the time, which lacked sufficient memory and processing speed to achieve substantial practical results. One setback was encapsulated in Moravec's paradox: Computers excelled at tasks like theorem proving (a subfield of automated reasoning and mathematical logic) and geometry problem solving, which were considered difficult for humans, but struggled with tasks that are easy for humans to perform – motor and social skills – like facial recognition or navigating a room without collisions (Arora, 2023; Zador, 2019).

In the mid-1980s the intersection of statistical mechanics and learning theory came into focus. During this time statistical learning from examples took precedence over traditional logic and rule-based AI. This shift was marked by two influential papers: Valiant's (1984) "A Theory of the Learnable," which laid the groundwork for rigorous statistical learning in AI, and Hopfield's (1982) development of the neural network model for associative memory. Hopfield's work sparked the widespread application of concepts borrowed from spin glass theory to neural network models. A pivotal moment in this evolution was the calculation of memory capacity in the *Hopfield model* by Amit, Gutfreund, and Sompolinsky (1985), followed by subsequent research endeavors. A more focused and in-depth application to learning models emerged through the pioneering work of Elizabeth Gardner, who harnessed *the replica trick* (Gardner, 1987, 1988) to compute volumes within weight spaces for straightforward feed-forward neural networks, encompassing both supervised and unsupervised learning models. Concurrently, the "backpropagation" method was popularized as the most popular method of training neural networks (Rumelhart, Hinton, & Williams, 1986).

Despite these notable advancements, the initial optimism had set unrealistic expectations. When the promised results failed to materialize, funding for AI dwindled, leading to a period of decline.

4. KNOWLEDGE-BASED SYSTEMS AND THE RISE OF EXPERT SYSTEMS:

a. 1970s to 1980s: From the 1970s through the mid-1980s, the ongoing evolution of AI continued to advance along the two primary paths that had been present from the outset: the enhancement of reasoning processes and the refinement of expert knowledge representation (Brock, 2018). AI researchers develop knowledge-based systems, where explicit rules determine decision-making. As an example, starting in 1972 and concluding in 1980, the MYCIN project played a pivotal role in assisting physicians in diagnosing a range of infectious blood diseases based on collections of clinical symptoms. What set MYCIN apart from DENDRAL was its approach of maintaining a clear separation between the rules (comprising expert knowledge) and the inference engine responsible for applying these rules. Expert systems were "the new new thing" in the 1980s, daily applied in roughly two-thirds of the Fortune 500 companies (Press, 2020).

5. MACHINE LEARNING EMERGENCE:

a. 1980s: Knowledge acquisition eventually became automated, although it did not unfold in the manner initially envisioned. The shift starts from rule-based systems to algorithms that can learn from data, inspired by the concept of neural networks from the 1950s. "Explanation Based Learning (EBL)" is introduced, a concept wherein a computer analyzes training data to formulate general rules, enabling the elimination of less essential data (DeJong & Mooney, 1986). NetTalk, a system that learns word pronunciation in a manner resembling a child's natural learning process, is also introduced (Sejnowski & Rosenberg, 1988).

b. 1990s: Decision trees, support vector machines, and other ML algorithms become prominent. A pivotal element in AI is the decision tree, an influential and adaptable instrument that has been instrumental in advancing the realms of machine learning and data mining methodologies. Although decision trees can be traced to the 1950s and 1960, in the 1990s ensemble methods came to the fore, uniting numerous decision trees to enhance the overall performance and precision of models (Frąckiewicz, 2023). The support vector machine (SVM), developed in 1963 and further refined by Vapnik in the 1990s, is another classification tool. These are powerful yet flexible supervised machine learning algorithms that fundamentally represent various classes within a multidimensional space through a hyperplane. SVM iteratively generates this hyperplane to minimize errors. The primary objective of SVM is to partition datasets into classes while seeking a maximum margin hyperplane (MMH) (Bandgar, 2021).

6. THE INTERNET ERA AND BIG DATA:

Late 1990s to 2000s: The rise of the internet provides a vast amount of data. Machine learning, especially through data-mining techniques, gains importance in businesses and technology sectors. Tim Berners-Lee and Robert Cailliau pioneered the World Wide Web and

created fundamental technologies like HTML, URLs, and HTTP during their tenure at CERN. This groundbreaking achievement marked the dawn of the internet age, enabling widespread and convenient access to information (Phillips, 2021). By 1996 digital data storage had become more cost-effective than storing information on paper (Morris & Truskowski, 2003). The domain google.com was registered a year later, in 1997, and subsequently launched in 1998. This marked the commencement of the search engine's ascent towards data dominance, as well as the initiation of various other technological advancements, particularly in fields like machine learning, big data, and analytics (Firican). More on machine learning below.

7. DEEP LEARNING REVOLUTION:

a. 2010s: Advancements in neural networks lead to the development of deep learning. In the year 2006, Hinton and colleagues introduced "Deep Learning." Deep Learning constitutes a subset within the realm of machine learning, focusing on algorithms influenced by the intricate workings of the human brain, known as artificial neural networks. These neural networks play a pivotal role in tasks like image classification, speech recognition, object detection, and content description. The emergence of deep learning marked a pivotal moment, leading to a resurgence in neural network exploration, often dubbed as the "new-generation neural networks." This resurgence stems from the remarkable accomplishments achieved by deep networks in effectively addressing a diverse range of classification and regression tasks (see e.g., Fong et al., 2021 on how deep learning convolutional neural network can be used to predict emotional responses to music).

b. Several important tools and platforms emerged in the 2010s to make AI development more accessible. Google's deep learning framework TensorFlow and Facebook's PyTorch (an open-source ML library based on the Torcj library) became prominent open-source deep learning frameworks, simplifying the creation of AI models and giving everyone the tool to build extraordinary models (Tyson, 2022).

Keras, integrated into TensorFlow, offered a high-level API for rapid prototyping. Scikit-Learn remained valuable for data analysis and machine learning (Géron, 2022). Jupyter Notebooks and Google Colab provided interactive environments for code, data, and visualization. AutoML tools like Google's AutoML and H2O.ai's Driverless AI automated aspects of model development. Microsoft Azure ML offered cloud-based machine learning services, while IBM Watson provided AI capabilities for businesses. These tools collectively empowered researchers, data scientists, and developers to engage in AI development and research.

8. MODERN ERA AND AI EXPANSION:

a. Late 2010s to 2020s: AI sees applications in almost every domain, from healthcare (Bohr & Memarzadeh, 2020) to finance (Goodell et al., 2021) to entertainment (Mateas, 2002) and education (Farrokhnia et al., 2023; Liang et al., 2023). The role of AI in data analytics (Kibria et al., 2018), natural language processing (Deng & Liu, 2018), and robotics (Bogue, 2014) becomes especially pronounced.

b. 2020s and Beyond: The convergence of AI with other technologies like quantum computing (Ajagekar & You, 2022; Choi et al., 2020) and augmented reality. Artificial intelligence and

quantum computing share several common features. Quantum computing, in particular, offers artificial intelligence and machine learning algorithms a significant advantage in terms of training speed and computational power, often at a lower cost (Abdelgaber & Nikolopulos, 2020). At the same time, the convergence of quantum computing and AI necessitates the achievement of various milestones to unlock their full potential in the realm of quantum computing AI (Rawat et al., 2022).

Augmented Reality (AR) entails enhancing our perception of reality by superimposing digital information onto objects viewed through a device. The fusion of AR and AI stands out as a prominent and imminent direction, acknowledged by numerous industries and academic circles. By harnessing the capabilities of AI, there exists significant potential for industries to enhance production speed, workforce training, manufacturing processes, error handling, assembly tasks, and packaging procedures. As is the case with quantum computing, there are numerous technological obstacles when applying AI to AR. Overcoming these, however, are beneficial, because with the infusion of AI capabilities, AR systems will gain the autonomy to operate within production environments with minimal human intervention. This synergy between AI and AR empowers these systems to seamlessly interact with both the manufacturing environment and human operators, thereby fostering the realization of Industry 4.0 (Devagiri et al., 2022).

Increased focus on general AI (artificial general intelligence – AGI). General AI refers to a form of artificial intelligence that possesses the ability to understand, learn, and apply knowledge across a wide range of tasks and domains, much like a human being. Unlike narrow or specialized AI, which is designed to excel at specific tasks or within predefined domains, general AI aims to exhibit human-like cognitive abilities and adaptability (Grudin et al., 2019; Pei et al., 2019). There are those that argue that there exists a genuine potential for substantial and interconnected progress in AGI design, engineering, evaluation, and theory in the relatively near future, possibly within the next few decades and potentially even sooner (Goertzel, 2014).

c. In the 2020s and beyond, chat AI has transformed the landscape of digital communication and information processing. These advanced AI systems, exemplified by models like GPT-4 from OpenAI, have become integral in various fields, providing accurate, context-aware, and often creative responses (see Ray, 2023 and Sohail et al., 2023 for comprehensive reviews). They assist in language translation, content creation, education, and even complex problem-solving. With continual advancements, these AI models are expected to become more sophisticated, offering even more personalized and nuanced interactions (Hassani & Silva, 2023). Their impact on industries such as customer service, healthcare, and education is profound, automating tasks and providing insights that were previously impossible. As we move forward, the integration of AI in daily life promises to be more seamless and impactful, heralding a new era of human-AI collaboration.

CHAPTER 2: BASICS OF AI AND MACHINE LEARNING

1. WHAT IS AI?

AI refers to the simulation of human intelligence in machines that are programmed to think and act like humans. The term can also be applied to any machine that exhibits traits associated with a human mind, such as *learning* through recognition of patterns or features in data, *reasoning* or determining future actions, *self-correction* through adjusting and refining of methods to improve results, and *problem-solving* (including methods like search and mathematical optimization). Early commercial movers such as IBM's Watson, simply referred to it as "software that learns." AI systems can also be trained to *perceive* their environment by recognizing voices, texts, and images. Technologies like computer vision and natural language processing (NLP) fall under this category. In essence, AI is about creating algorithms that allow computers to perform tasks that would ordinarily require human intelligence. It spans a wide range of applications, from simple calculators that solve math problems to sophisticated systems that can understand language, recognize patterns, and make decisions (Spatharioti et al., 2023).

The question of how to define AI is key in this respect. Beginning with Turing's question "Can machines think?" (1950), through McCarthy et al. (1955) idea that "the artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so behaving," to Floridi & Cowls (2019) argument that "the classic definition enables one to conceptualize AI as a growing resource of interactive, autonomous, and often self-learning agency that can deal with tasks that would otherwise require human intelligence and intervention to be performed successfully" definitions mostly focus on the link to human behavior and how AI might complement or diverge from it[1]. In contrast, Tegmark (2018) defines AI simply as "non-biological intelligence." Building on Russell's (2020) observation that "machines are intelligent to the extent that their actions can be expected to reach their objectives," Cooke & Passingham (2022) consider it essential to clearly specify the AI program's objectives, how it reaches conclusions, and the ethical and moral guidelines directing its operation, ensuring transparency and compliance in a global context.

AIs impact on businesses, economies, and societal structures is transformational: elevating efficiency and automation, boosting analysis and insights through big data processing (Agbehadji et al., 2020; Batko & Slezak, 2022; Bragazzi et al., 2020) and predictive analysis (Ali & Djalilian, 2023; Choi et al., 2023; Khan, 2023; Kumar et al., 2022; Lim et al., 2023; Perkins, 2023; Tong & Zhang, 2023), increasing user experience through personalization (customized user experiences) and almost-natural interactions (chatbots and voice assistants), fostering innovation in products and services through new offerings and service enhancement, supporting decision-making through complex problem solving and real-time decisions (for e.g., in sectors like finance and logistics), providing scalability by handling growth and globalizing operations, boosting creativity and design through content creation (Gozalo-Brizuela & Garrido-Merchan, 2023; Ipsos, 2023) and optimizing design (Brossard et al., 2020).

[1] See also IAB AI Standards Working Group (2021) definition: Artificial intelligence or AI is the empowerment of machines to use reason and understanding to complete tasks, unlike natural intelligence, which humans and animals employ and involve conscious reasoning and understanding (p. 34).

AI is used to address societal challenges in healthcare and environment, as well as to detect bias and enhance transparency, thus providing a more ethical and responsible notion of AI (Ipsos, 2023).

2. MACHINE LEARNING AND DEEP LEARNING: AN OVERVIEW

Machine Learning (ML):

In recent years, Artificial Intelligence (AI), especially Machine Learning (ML), has experienced rapid growth in the realm of data analysis and computing. This evolution enables applications to operate intelligently. ML empowers systems to autonomously learn and improve from experience without the need for explicit programming. In the *Fourth Industrial Revolution* (4IR or Industry 4.0), namely the ongoing automation of traditional manufacturing and industrial practices, ML stands as one of the most prominent technologies. Instead of being explicitly programmed to perform a task, a machine learning system learns from data, thereby improving its performance over time. ML algorithms are thus considered the cornerstone for intelligently processing this data and creating real-world applications (Sarker, 2021a). These learning algorithms can be categorized into four primary types: supervised, unsupervised, semi-supervised, and reinforcement learning (see below). The popularity of these learning approaches has been steadily rising, as demonstrated in Figure 2.

Figure 2: Data collected from Google Trends (Sarker 2021a: 159)

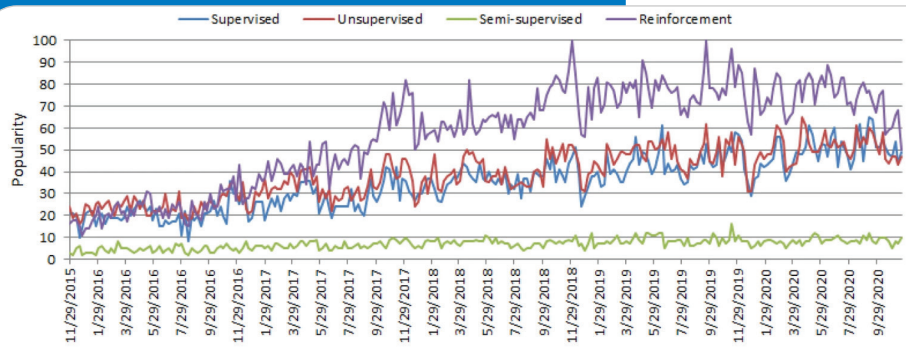

The worldwide popularity score of various types of ML algorithms (supervised, unsupervised, semi-supervised, and reinforcement) in a range of 0 (min) to 100 (max) over time where x-axis represents the timestamp information and y-axis represents the corresponding score

Here's a breakdown:

1. Types of Machine Learning (see Figure 3):

a. Supervised Learning: Supervised learning involves the machine learning process of acquiring knowledge from labeled input-output pairs to establish a mapping function. It relies on annotated training data and a dataset of training instances to deduce this function. Supervised learning is employed when specific objectives need to be achieved using a defined set of inputs, essentially adopting a task-oriented approach. The most prevalent tasks within supervised learning encompass "classification," which segregates data, and "regression," which models the data (Cunningham et al., 2008). For instance, predicting the class label or sentiment of textual content, such as a tweet or a product review, known as text classification, serves as a prime illustration of supervised learning.

b. Unsupervised Learning: Unsupervised learning delves into the analysis of unlabeled datasets autonomously, operating without human intervention—thus embodying a data-driven approach. It serves various purposes, including extracting generative features, uncovering meaningful patterns and structures, discovering inherent groupings within data, and facilitating exploratory endeavors. Among the prevalent tasks within unsupervised learning are clustering, density estimation, feature learning, dimensionality reduction, the discovery of association rules, and anomaly detection, to name a few (Ghahramani, 2003).

c. Semi-supervised Learning: Semi-supervised learning can be described as a fusion of the supervised and unsupervised approaches since it operates on both labeled and unlabeled datasets. Consequently, it occupies a middle ground between learning "without supervision" and learning "with supervision." In real-world scenarios, labeled data can be scarce in numerous contexts, while unlabeled data abound, making semi-supervised learning particularly valuable (Zhu, 2005). The primary objective of a semi-supervised learning model is to deliver improved prediction outcomes compared to those achievable using only labeled data within the model. Semi-supervised learning finds application in various domains, including machine translation, fraud detection, data labeling, and text classification.

d. Reinforcement Learning: Reinforcement learning is a category of machine learning algorithms that empowers software agents and machines to autonomously assess the most effective behavior within a specific context or environment to enhance their performance. This approach is fundamentally driven by the characteristics of the environment itself. The model learns by interacting with an environment and receiving feedback (rewards or punishments) based on its actions, with the ultimate objective of utilizing insights gained from interactions with the environment to make choices that maximize rewards or minimize risks (Kaelbling et al., 1996). Reinforcement learning serves as a potent tool for training AI models, particularly in domains where automation enhancement and operational efficiency optimization are paramount. Such domains include robotics, autonomous driving systems, manufacturing processes, and supply chain logistics. However, it is not typically the preferred choice for addressing simple or straightforward problems.

Figure 3: Machine Learning Types (Sarker 2021a: 160)

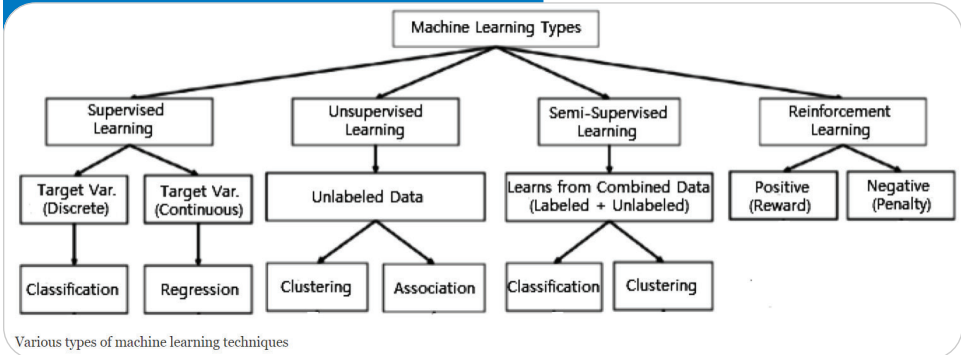

Various types of machine learning techniques

For a detailed review of the main ML techniques, see here.

2. Various Applications:

- Predictive analytics (e.g., stock price prediction)

Predictive analytics powered by machine learning facilitates intelligent decision-making through data-driven insights (Cao, 2017; Mahdavinejad et al., 2018). It uses relationships between variables from past data to forecast outcomes [Han et al., 2009]. AI models can predict future trends based on historical data, allowing industries to preemptively adjust strategies. In education (Choi et al., 2023; Lim et al., 2023), disease diagnosis (Khan, 2023; Kumar et al., 2022), specific research fields such as synthetic biology (Tong & Zhang, 2023), not to mention writing in academia (Perkins, 2023) and scientific fields (Ali & Djalilian, 2023), researchers are predicting how AI will affect various areas and using AI to predict research trends. Accurate predictions benefit various sectors, from government to healthcare, banking, e-commerce, and more. In the context of advertising, predictive analytics leverages AI and machine learning to target audiences and optimize ad placements, focusing on personalization and real-time content delivery. Image, speech recognition, and pattern recognition.

- Image recognition

Image recognition through machine learning (Fujiyoshi et al., 2019) is highly prevalent in real-world applications. Through image recognition one can identify objects in digital images. Labeling x-rays, character and face detection, and social media tagging are all examples of image recognition. Speech recognition (Chiu et al., 2018), used in platforms like Google Assistant and Alexa (López et al., 2017), leverages sound and linguistic models through machine learning. Pattern recognition (Anzai, 2012) automates the identification of data patterns using techniques like classification, clustering, and sequence labeling. In advertising, machine learning-enhanced image and speech recognition can be transformative, enabling advanced targeting and personalization. For instance, image recognition can be used for analyzing consumer behavior through social media imagery, while speech recognition powers voice-activated advertisements, offering interactive and tailored consumer experiences.

- Medical diagnosis

Machine learning aids in various medical tasks, from disease prediction to patient management (Fatima et al., 2017; Nilashi et al., 2017). For instance, in the case of Covid-19, learning techniques have been instrumental in classifying high-risk patients, predicting outbreaks, and understanding the virus's origin (WHO; Kushwaha et al., 2020). Machine learning forecasted potential COVID-19 spread, guiding regions to prepare. Additionally, deep learning offers solutions in medical image processing. In advertising, the application of machine learning techniques used in medical fields, such as disease prediction and patient management, can be adapted for consumer health trends analysis and targeted health-related campaigns. For example, deep learning in medical image processing can inspire advanced visual analytics in advertising, enhancing ad personalization and effectiveness.

- Recommender systems (e.g., movie recommendations)

Machine learning powers product recommendations, a hallmark of modern e-commerce. It analyzes consumers' buying patterns and preferences to personalize suggestions. By studying browsing trends and click-through rates, e-commerce platforms, like Amazon (Marchand & Marx, 2020), optimize inventory, logistics, and prevent stockouts. Machine learning enhances the shopping experience through tailored content and helps retain and attract customers.

- NLP and sentiment analysis

NLP enables computers to interpret and understand language (Otter et al., 2020). It powers applications like virtual assistants, chatbots, speech recognition, and machine translation. Sentiment Analysis, a subset of NLP (Ravi & Ravi, 2010), determines public opinion from sources like blogs, reviews, and social media. It helps businesses gauge brand or product sentiments, categorizing feedback as "positive", "negative", "neutral", or other specific emotions. In essence, machine learning assists in extracting and categorizing these sentiments from text. In advertising, NLP and Sentiment Analysis enable advertisers to extract and categorize public sentiments from online sources like blogs and social media, providing insights into brand perception and guiding the creation of more resonant and effective advertising strategies. This application of NLP in analyzing feedback aids in tailoring advertisements to align better with consumer attitudes and trends.

CHAPTER 3: INTEGRATING AI INTO MARKETING AND ADVERTISING RESEARCH – A COMPREHENSIVE REVIEW OF THE LITERATURE

Integrating AI into marketing and advertising research involves leveraging machine learning algorithms, data analytics, and computational tools to gain insights into consumption, consumer studies, and consumer research including consumer behavior, how to optimize marketing campaigns, and enhance the personalization of advertising content. Brand et al. (2023) for instance, explore the potential of Large Language Models (LLMs) in market research applications, focusing on GPT3.5 ability to mimic consumer preferences and behavior. Reisenbichler et al. (2023) explore the integration of LLMs in creating advertising content for search engines, highlighting the development of an application layer over OpenAI's GPT models to generate ad text tailored for search engine advertising and predict advertising costs. Similarly, Spatharioti et al. (2023) explore the impact of LLMs on consumer product research, particularly in the context of search engine use. Hayes et al. (2021) compared graduate students' sentiment analysis of social media posts on Nike's "Dream Crazy" ad featuring Colin Kaepernick to that of two major tools, Crimson Hexagon and LIWC, finding that these tools' accuracy in categorizing sentiments as positive, neutral, or negative was only about one-third, equivalent to random chance.

Four academic literature reviews on the topic of the increasing relevance of *AI in marketing* are noteworthy. We present each in brief. A more extensive review can be found here.

Vlačić et al. (2021) review 164 articles, identifying AI's impact on marketing and advertising research across four overarching themes:

1. Marketing Channels: AI enhances efficiency in marketing channels (Bock et al., 2020; Wirtz et al., 2018), with applications in predicting customer preferences and improving product distribution, as seen with North Face and Amazon (Sjödin et al., 2018).

2. Marketing Strategy: AI reshapes strategic marketing, balancing massification and customization (Du et al., 2003), and integrating luxury brands with mass markets (Kumar et al., 2020c; Paul, 2019).
3. Performance: AI's predictive capabilities are superior for performance evaluation (Bock et al., 2020; Russell & Norvig, 2016; Syam & Sharma, 2018) and impact competitive advantage and customer value creation (Paschen et al., 2020; Riikkinen et al., 2018; Diaz, 2017).
4. Segmentation, Targeting, and Positioning (STP): Advances in STP utilize AI for customer base management, leveraging neural networks for market segmentation (Fish et al., 1995; 2004; Ha et al., 2005) and customer profiling (Lei & Moon, 2015; Wu et al. 2015; Belk, 2016; Pitt et al., 2018).

Mustak et al. (2021) provide an in-depth analysis of AI in marketing, identifying ten dominant research topics divided into consumer and organization/strategy-related themes. Consumer research includes AI in sentiment analysis (Chong et al., 2016; Zhang et al., 2018), customer satisfaction (Ansari & Riasi, 2016; Baumann et al., 2012), eWOM (Pantano et al., 2019), brand management (Haryanto et al., 2015), customer loyalty (Ballestar et al., 2019), and relationship management (Bejou et al., 1996). Organizational themes cover AI in B2B automation (Cascio et al., 2010), market performance (Erevelles et al., 2016), service innovation (Yu, 2020; Van et al., 2019), and strategic marketing (Lin & Kunnathur, 2019; Netzer et al., 2012).

Mariani et al. (2021) conducted a systematic review of AI research across marketing, consumer research, and psychology, identifying eight key thematic areas in AI research: (1) memory and computational logic; (2) decision making and cognitive processes; (3) neural networks; (4) machine learning and linguistic analysis; (5) social media and text mining; (6) social media content analytics; (7) technology acceptance and adoption; and (8) big data and robots. This work showcases the diverse and expanding scope of AI's application in these fields.

Notably, all three literature reviews demonstrate a rapid and exponential growth in the evolution of academic publications on the topic of AI in marketing over time, with Vlačić et al. (2021) pointing to a significant spike in 2017 onwards.

Finally, Haleem et al. (2022) analyzed 217 academic publications, highlighting AI's diverse applications in marketing segments like personalized marketing, predictive analytics, customer segmentation, automated decision-making, and customer engagement.

In industry research, AI in market research covers a range of topics. From how AI will impact jobs (Qualtrics 2018 surveys industry perspectives) through AI's potential to both enhance and disrupt marketing and sales (Deveau et al., 2023; see also Ruden, 2023 on how AI can optimize paid search campaigns) and to practical guides designed to help those already working with AI or looking to leverage it in their businesses (IAB AI Standards Working Group, 2021).

Ostler & Kalidas (2023) emphasize the potential of Large Language Models for efficiency and data analysis, predicting the use of LLMs in three primary use cases. Firstly, to enhance efficiency by automating tasks, such as the manual coding of open-ended responses. Secondly, to improve capabilities in areas like analyzing vast amounts of data. Lastly, LLMs create new opportunities, such as generating multiple versions of a concept and using

another AI system to evaluate each, thereby identifying the most effective option. This mirrors Analytics Partners (2023)'s view on Generative AI's capabilities and the need for substantial human involvement, oversight, and review of AI outputs, without which AI cannot effectively develop reliable and precise commercial analytics models.

According to Manole (2023) the market research industry is witnessing a transformative shift with the integration of AI. First, AI-Powered Market Research for Price Optimization is revolutionizing how businesses determine pricing strategies by meticulously analyzing market data to balance value and profitability. Simultaneously, Emotion AI is being employed in Customer Experience (CX) Market Research, enabling companies to identify and enhance key customer journey touchpoints through emotional and feedback analysis. The Rise of AI-Powered Survey Technology is another significant trend, where AI is streamlining the processing of free-text survey responses, transforming them into comprehensive and actionable insights. Additionally, Inclusivity through AI Models is gaining momentum, focusing on understanding and catering to a diverse range of customer demographics, thereby personalizing business strategies. Lastly, AI's application in Social Media Market Research is proving invaluable for analyzing user behavior and engagement, as well as predicting content performance, thereby refining targeted marketing strategies.

Moriarty (2023) identifies three prominent ways AI can enhance marketing: First, by analyzing digital trends and consumer sentiments to inform market research and strategy development. Second, through predictive analytics for efficient and insightful creative testing (see also Ho et al., 2022 discussing how machine learning (ML) can predict new product 4innovation success). Finally, AI optimizes campaign performance by providing real-time insights, enabling rapid adjustments and improvements.

Huff & Bonde (2022) discuss how the emerging field of AI-enabled consumer intelligence (AICI) is reshaping how businesses understand and interact with consumers. Since these platforms are designed to enable enterprises to gather insights from various data sources, both external (like social media, web data, and consumer data) and internal (such as CRM and website data), they can help brands grapple with the volume of consumer-generated data and uncover predictive and prescriptive insights.

Deveau et al. (2023) highlight how gen AI is revolutionizing marketing and sales, especially in B2B and B2C sectors, emphasizing the potential of gen AI in areas like customer experience, productivity, and growth. They discuss the automation of a significant portion of sales functions, the importance of advanced sales technology and hyper-personalization, and the impact of AI on various aspects of the customer journey. In contrast, Kerwin (2023) reflects on the cautious or wary approach with which marketers address Gen AI, underlining the importance of addressing crucial aspects including formulating a clear strategy, ensuring adequate training for usage, allocating the necessary budget, and understanding the legal implications involved before incorporating the outputs of AI tools. Relatedly, Hardcastle (2023) discusses the unintended consequences of Gen AI content in marketing, including the inadvertent production of harmful content, the impact of inconsistent brand tone of voice, the environmental impact of generating absurd quantities of content, and how an overreliance on Gen AI tools can stifle creativity.

Alongside the possible benefits of integrating AI into marketing, there are also voices calling to consider the inherent challenges in AI's multifaceted influence. For instance, Campbell (2023) critiques the limited scope of existing AI application frameworks in market research citing *The Cross-Industry Standard Process for Data Mining (CRISP-DM)* as an example. He calls for a voluntary agreement among market research companies and practitioners to guide the adoption of these technologies and ensure they are deployed ethically and responsibly. Cooke & Passingham (2022) discuss the need to shape ethical frameworks to retain the industry's ability to self-regulate (see more in chapter 5 on ethics).

AI advertising, namely, "brand communication that uses a range of machine functions that learn to carryout tasks with intent to persuade with input by humans, machines, or both" (Rodgers, 2021, 2)[2] carries with it much promise (and perhaps also peril). Huh et al. (2023) identify four domains of advertising practice and research that are undergoing profound transformations due to the influence of AI, creating new avenues for exploration, and necessitating further research efforts (see Figure 4):

(1) *Consumers' experience of advertising*. AI is revolutionizing advertising by enabling novel human-AI interactions with virtual chatbots, influencers, and brand spokespersons. This transformation prompts inquiries into how consumers engage with AI persuasion agents, distinguish them from humans, and the societal implications of this AI-powered world. Additionally, generative AI technology is reshaping search advertising, altering search behaviors and result presentation. Lastly, the emergence of AI-powered virtual influencers raises questions about their impact on social influencer marketing and whether existing advertising theories remain relevant in this evolving landscape.

(2) *Societal and policy implications related to the truthfulness of AI-generated content*. The rise of generative AI technology raises concerns about the spread of fake information in AI-generated content. Researchers can address this through advertising literacy training and considering regulations for ethical AI-powered advertising.

(3) *The analytical considerations of data and algorithms*. AI advertising poses ethical concerns regarding data collection, transparency, and bias. Research is needed on evolving data privacy issues, covert data capture's implications, and consumer awareness. Opacity in advanced AI models hampers transparency, affecting ethical advertising. Algorithmic bias in data-driven advertising and its societal impact also require examination.

(4) *The functioning of the advertising industry*. AI's impact on the economy, especially job displacement, is a long-debated issue. While some jobs have been replaced by automation and AI, the advertising industry has been relatively insulated. However, the rise of generative AI raises questions about the future of creativity in advertising and which tasks will still require human intervention, impacting revenue models, organizational structures, and client-agency relationships. As an example, consider Fond et al. (2021) use of deep learning models to improve emotion-based advertising in digital media.

Ford et al. (2023) conducted a comprehensive review aiming to map the field's evolution. Analyzing *AI advertising* academic articles published between 1990 and 2022, they identify a publication trend wherein despite the infancy of AI advertising research, since 2018 there

[2] However, definitions of "artificial intelligence and advertising" vary widely, see Rodgers (2021) for elaboration.

Figure 4: Illustration of AI advertising research areas (Huh et al., 2023: 479)

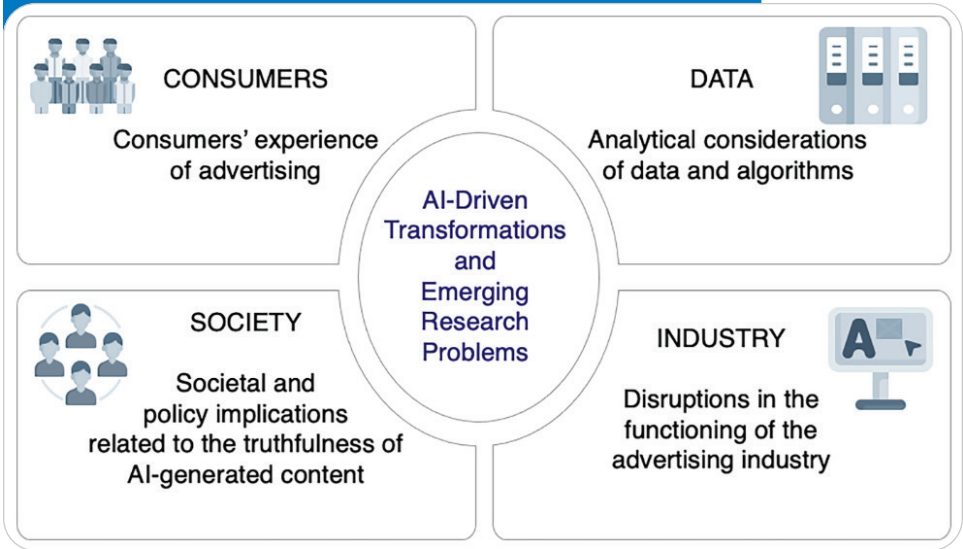

is enormous growth in academic articles related to AI advertising (see Figure 5). Analyzing 72 articles they identify four clusters as key focus areas of AI advertising research: AI-driven advertising innovation, Computational advertising (CA), Programmatic advertising (PA), Ad Effectiveness in AI Advertising.

Figure 5: AI advertising publication trend (Ford et al., 2023: 4)

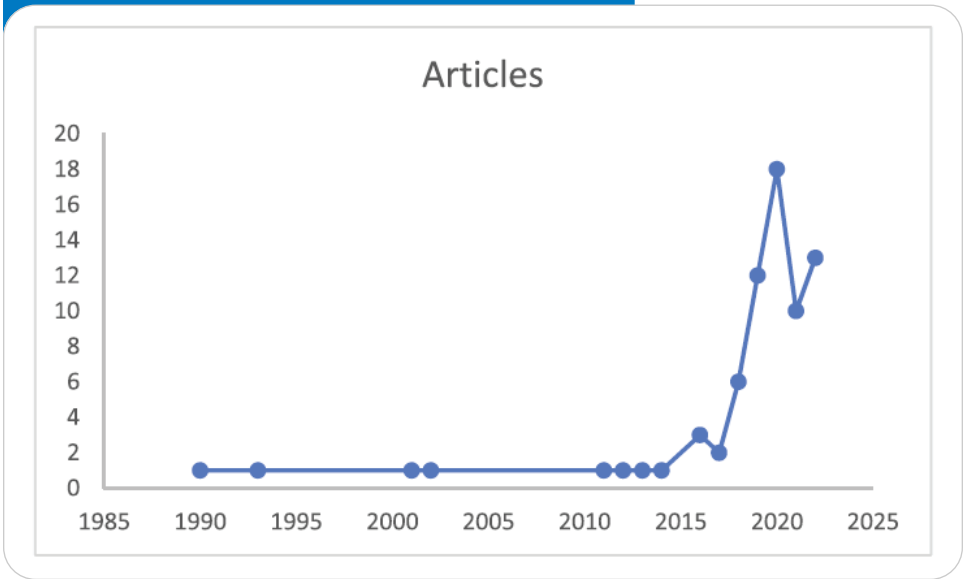

Cluster 1: AI-driven advertising innovation features 18 articles. More than half are conceptual, discussing AI's role in advertising and suggesting future research avenues (Kietzmann et al., 2018; Li, 2019; Coffin, 2022). These works, along with others, examine AI applications in various areas. For instance, Qin & Jiang (2019) focus on smart advertising, namely how AI is used in various stages in the advertising process such as consumer insight discovery, ad creation, media planning and buying and ad impact evaluation. Campbell et al. (2022) examine manipulated advertising, specifically how consumers respond to 'synthetic ads' that

rely on deepfakes and generative adversarial networks (GANs). Bakpayev et al. (2020) focus on programmatic creative to show how AI ads are effective in rational appeal, yet they lack in emotional aspects compared to human-made ads. Deng et al. (2019) and Vakratas & Wang (2020) explore enhancing ad creativity through smart ads.

Cluster 2: Computational advertising features 19 articles. Articles in this cluster focus on how AI techniques like predictive modeling, programmatic advertising, algorithms, and machine learning are leveraged to enhance advertising efficiency, and explore how these technologies use data for tailored ads and targeting. For instance, Malthouse et al. (2019) merge programmatic advertising with recommender systems, creating content delivery algorithms. Similar algorithms for consumer profiling are discussed by Neumann et al. (2019), with applications in traditional and mobile advertising (Li & Du, 2012; Guitart et al., 2021). Chen et al. (2019) advance programmatic creative for automated ad creation, while Liu-Thompkins et al. (2020) and van Noort et al. (2020) focus on automating brand content. Araujo et al. (2020) examine brand-consumer experiences in this context. Zimand-Sheiner & Earon (2019) investigate the impacts of computational techniques on account planning through a qualitative study with managers.

Cluster 3: Programmatic advertising features 14 articles focusing on the increasingly popular programmatic advertising within AI-driven computational advertising. This automation of media buying enhances ad placement efficiency and performance. Samuel et al. (2021) explore its processes and consumer concerns related to personalization. Palos-Sanchez et al. (2019) address privacy concerns, finding no significant increase over time. Other research includes methodology development for search and display advertising optimization (Gong et al., 2017; Miralles-Pechuán et al., 2018), emphasizing ad relevance (Kononova et al., 2020) and consumer motivations (Lee & Cho, 2020). This cluster primarily examines website advertising in the context of programmatic advertising's impact and consumer response.

Cluster 4: Ad effectiveness in AI advertising features four articles. Shumanov et al. (2022) conduct a mixed-method study to evaluate the effectiveness of AI-predicted personality-based ads, showing that for most personality types matching consumer personality with congruent advertising can lead to more effective consumer persuasion. Matz et al. (2019) assess image appeal in advertising computationally. Huh et al. (2020) explore viral ad diffusion, finding that high source trust among consumers enhances engagement with viral ads. Lastly, Roy et al. (2017) develop a Trust Score Algorithm for social media advertising, demonstrating that both advertiser and sender trust positively influence ad effectiveness.

In industry research, research focuses on how AI presents opportunities and challenges to advertisers and their agencies (e.g., Gralpois, 2023; Hsu & Lu, 2023; Vranica, 2023). Forrester's "B2C Marketing CMO Pulse Survey" predicts AI will enhance creative agencies, potentially close digital shops, and trigger more reviews in the coming year. The survey anticipates the leading 10 agencies to invest around $50 million collectively in partnerships for developing bespoke AI solutions for clients (Stam, 2023). Erdem & Sidlova (2023) (see also Ostler et al., 2023) tested some fully AI-generated ads along those partially developed by GenAI (like script writing or image creation), using Kantar's AI-powered ad-testing tool that leverages data from 250,000 real-world ad tests. They show how GenAI ads performed strongly, but quality was variable, thus arguing for the role of AI as partner that helps agencies and marketers explore new ways to be creative.

When delving into how AI is used for *advertising research*, however, research is scarce. Indeed, two call for papers have been recently issued on this topic: the Journal of Advertising Research, on "*How Can Advertisers Leverage Ai And Generative AI*?" and the International Journal of Advertising on "*Evolution, Challenges, and Opportunities of AI-Generated Advertising."* As a first step to address this lacuna, in what follows, we present findings from 8 case studies, each showcasing a different aspect of advertising research and how AI can be used (with more or less success) to address these.

CHAPTER 4: ARF CASE STUDIES – AI IMPLEMENTATIONS IN ADVERTISING RESEARCH

This chapter presents a series of case studies that showcase the diverse applications of AI in the field of advertising research. Each case study is presented in a condensed format, providing key insights and findings in a succinct manner. For readers interested in exploring these topics in greater depth, the chapter oftentimes offers longer versions of the case studies that include the wide range of prompts the research team used and the various responses provided by the AI platforms. The additional material can be accessed in the link provided at the end of each case study.

Also, we recognize that AI has had extensive use in prgrammatic targeting and creative testing. Consequently, they are not considered among these case studies. However, The ARF is currently conducting a research initiative evaluating several commercial offerings that use AI in creative testing and we will be reporting these results as the project develops.

1. CASE STUDY 1: DATA INGESTION

ChatGPT 4

By using ChatGPT 4's Advanced Data Analysis Beta Feature, the user can upload a variety of files in different formats (including CSV) that can be parsed by the model. This makes ChatGPT 4 the only model that can ingest full datasets, giving it full access to the information one would need to perform the data analysis.

At the same time, some amount of pre-processing needs to be done before the files are uploaded to ChatGPT 4. For example, oftentimes ChatGPT 4 has difficulty extracting information from WORD and PDF format files that include text and charts, images, or graphs (see below Erwin Ephron Case Study for further elaboration). One way to solve this is to identify headers, paragraphs, and general structure of document; another way is to parse files to more manageable data. Finally, if nothing else works, a simple 'copy' and 'paste' will do.

ChatGPT 4 also demonstrated difficulty understanding datasets downloaded directly from Qualtrics. In the Privacy Report survey, for instance, even after specifying that there are 3 header rows and that the second row contains the actual question, ChatGPT 4 was unable to identify data. Thus, it was easiest to delete all header rows except the second one that contains the question, and re-upload the data set.

Bard

Bard claims to be able to access publicly available documents via links yet this is not the case. Consequently, it is easier to resort to inputting the data in raw text format.

Bard also purports to possess an expedited mechanism that facilitates direct linkage with files stored within Google Drive or Drop Box, inclusive of CSV files. However, when uploading data files to Google Drive and asking Bard to access them, the outputs are not aligned with the documents' actual content. Additionally, Bard's attempts to analyze the unprocessed data file from Google Drive resulted in Python code for statistics not present in the dataset or previous reports.

Claude AI

Claude AI boasts the most streamlined data ingestion process, permitting file uploads including PDF, txt, CSV, albeit with a size limit that prevents the upload of the full data file (5 documents max, 10MB each).

Claude AI did ingest the original files in CSV format but said that it could not reproduce any of the data within the file due to copyright constraints. As a result, it returned generalized code in Python that one might use to perform a data analysis, which was not specific to the data file supplied (for example, all variable names and column names had been falsified). Thus, the code would require significant altering in order to be used for analysis on the actual data set.

With a free Claude AI account, user interactions with the model are capped. As a result, comprehensive analyses might need to be split up across multiple times/days to incrementally achieve the intended outcome.

Data ingestion capabilities	ChatGPT 4.0	Bard	Claude.AI
	Text input and file upload (including CSV)	Claims to be able to ingest text input and connection to files stored in Google Drive, yet this was unsubstantiated.	Text input and file upload (including CSV) but ingesting a CSV did not result in a useful output.

2. CASE STUDY 2: SURVEY DESIGN

This case study highlights the utilization of AI in crafting survey questions, emphasizing its role in streamlining the process and generating innovative ideas. Our findings show that AI for survey design is a repetitive but time-saving method. We'll demonstrate this through examples, specifically in the context of media research surveys. While our primary focus is on Advertising Research, we've also employed AI, particularly Large Language Models (LLMs), in creating Awareness and Usage surveys. This includes integrating prominent brands into the surveys and adapting them for different regions, like modifying a survey for Mexico by substituting U.S. brands with leading Mexican ones. The findings of this case study underscore

the utility of using AI for designing survey questions yet highlight the need for iterative prompting and refinement in AI-assisted tasks.

Our initial prompt for all three AI platforms was as follows:

> I would like you to help me design a survey using best practice in question design. This survey should explore American adults' usage of media devices, how they get their television programming, their usage of streaming services and which services do they use or subscribe to. Can you start by suggesting an outline for the dimensions that the survey should cover. Do this from the perspective of a professional media researcher.

All three models produced reasonable outlines for the survey generally starting off with demographic questions, then moving to television services and then streaming services.

ChatGPT 4

Initially, ChatGPT 4 included 'Household Income' in 'Basic Demographic Information' section but moved this to end of survey upon request. It also neglected to include 'Race and Ethnicity,' and added this once requested.

ChatGPT 4 was effective in formatting questions similar to a survey authoring system, showing strong capabilities in handling race and region questions. Although at first, ChatGPT 4 combined 'Race' with 'Ethnicity,' it easily revised this upon prompt. While Bard provided a wider variety of choices for racial identity, Chat's use of a "check all that apply" format suggested a more nuanced understanding of multiracial identities.

In addition to providing an outline to the survey, ChatGPT 4 included 'Best Practices for Question Design,' underlining the importance of clarity, neutral wording, balanced answer options, and pilot testing, among other things.

ChatGPT 4's output was more easily prompted into a format that a programmer would create for an online survey. This model demonstrated adaptability in conditional question presentation, where questions dynamically changed or were suppressed based on previous responses.

To sum, while ChatGPT 4 required initial assistance with formatting, it demonstrated a rapid learning curve. Its capability to dynamically present device-related questions was notably superior compared to other LLMs. If the API were integrated into a system, the process could be further streamlined. By assigning names to specific question formats, the instructions for creating surveys could be greatly simplified, relying on these predefined formats for guidance.

Bard

Bard had issues with maintaining a vertical format for answers.

This model struggled with adapting to a 9-way region question, often defaulting to a 4-way region format.

It listed sources for its responses, although some were irrelevant or dead links. For instance, Bard provided 'Internships at Mystic Seaport Museum' as a link for choosing Race and Ethnicity categories, and www.freedoniagroup.com/industry-study/recreational-vehicles-in-the-us-by-product-6th-edition-3624.htm, which is a dead link.

Bard provided diverse options for race and ethnicity questions. For instance, for Race, it gave the following response:

Race (choose option based on your target audience and research goals):

- **Option 1 (US Census Bureau standard):**
  - White (including Hispanic or Latino)
  - Black or African American
  - American Indian or Alaska Native
  - Asian
  - Native Hawaiian or Other Pacific Islander
  - Some Other Race
  - Two or more races
  - Prefer not to say
- **Option 2 (Broader categories):**
  - Caucasian
  - African American/Black
  - Asian American
  - Native American/Indigenous
  - Other
  - Multiracial
  - Prefer not to say
- **Option 3 (Additional granularity, consult with survey experts for appropriateness):**
  - White (Non-Hispanic)
  - White (Hispanic)
  - Black or African American
  - Asian Indian
  - Chinese
  - Filipino
  - Japanese
  - Korean
  - Vietnamese
  - Other Asian
  - Native Hawaiian or Other Pacific Islander

- American Indian or Alaska Native (Tribal Affiliation)
- Some Other Race
- Two or more races
- Prefer not to say

Claude AI

Claude AI quickly adapted to changes in question format, such as switching from bullets to checkboxes, but had some spacing issues in its formatting.

It had a faster response to requests for a 9-way region question compared to Bard and showed decent performance in conditional presentation of devices.

Claude AI initially placed demographic questions at the end, which was later adjusted upon request. Similarly, it did not initially include education, race, and ethnicity in demographic.

It eventually got to a concise conditional presentation of devices but was not as elegant as ChatGPT 4's version.

TASK	ChatGPT	BARD	Claude
Outline	Provided good insight	Provided good insight	Provided good insight
Formatting	The best, reduced the conditional presentation similar to an authoring system	Frequently formatted categories horizontally rather than vertically	Frequently formatted categories horizontally rather than vertically
Handling Race	Anticipated multiracial: Check all that apply	Presented interesting options	Followed prompts
9-Way Region	Followed prompts	Couldn't give up on 4-way region	Followed prompts
Listed Sources	No	Yes, but some were irrelevant or dead links	No
Changing Bullets to Checkboxes	Had to tell Chat to use checkboxes in section 2 despite telling it to do so in section 1	Used checkboxes in section 2 once presented as a question	Maintained checkboxes in section 2 although asked
Formatting the conditioned grid, that is, only present the grid column if the panelists meet the condition	The best, most like a survey authoring system	Would have worked for a programmer	Would have worked for a programmer

Key Takeaways:

- Flexibility and Adaptability: ChatGPT 4 outperformed the other models in adapting question formats and conditional presentation, making it more suitable for drafting surveys that resemble those created by professional survey programmers.
- Handling issues of Race and Ethnicity: Bard offered interesting options for race and ethnicity but was less flexible in adapting to specific survey structures like the 9-way region format.
- Speed and Formatting: Claude AI showed quick adaptability in question formatting and adjusting to specific requests, such as the 9-way region question, though it was less elegant compared to ChatGPT 4's version.
- Learning and Improvement: The use of these LLMs in survey design demonstrated a learning curve, where the quality of output improved with more specific and refined prompts over time.
- Overall Effectiveness: All three AI systems represent effective methods for beginning a survey draft, with each having unique strengths and weaknesses in various aspects of survey design. We expect a frequent user of these LLMs would learn how to better prompt after programming a few surveys.

For a more detailed review of the prompts and LLM responses see.

3. CASE STUDY 3: SURVEY ANALYSIS

In this case study, we processed The ARF Privacy Study[3] survey data functions for frequency distributions, focusing on single-select and multiple-select questions. An ARF researcher, building on the 2021 and 2022 Privacy Studies which included topics like online privacy concerns and the customization of advertising through online behavior, then generated fresh insights for the 2023 survey. This involved exploring new questions not previously asked, such as perceptions of personalized advertising, and analyzing data across novel demographic segments, including education levels, race, and political affiliation, for the first time.

Next, we engaged AI tools – ChatGPT 4, Bard, and Claude AI – to assess their ability to interpret the survey data and we would then compare their analyses with human interpretations. The goal was to determine if these AI models could independently generate analysis based on the dataset and first training the LLMs on the Privacy Reports from 2021 and 2022.Since protected data should not be uploaded or shared with AI, a level of human pre-processing was required to remove any field that was not completely anonymized. For this survey, this included location data and IP address.

[3] The ARF Privacy Study is an annual survey conducted by the ARF that examines American perspectives on privacy, data collection (particularly in advertising), trust in institutions like medical, financial, and government bodies, and understanding of digital advertising terms. In its sixth year, 2023, the study gathered over 1,300 responses from U.S. participants, ensuring diversity in age, gender, and regional representation.

ChatGPT 4

The prompt:

> Attached is a data file that was collected in 2023, and 2 previous versions of reports that were written using similar data collected in previous years (2022 and 2021). Based on the 2021 and 2022 reports, could you please perform a similar analysis on the 2023 data and output the results?

The results:

ChatGPT 4 first segmented the privacy reports into distinct sections for focused analysis. ChatGPT 4 then extracted column headers from the dataset in batches of 50, requesting user input to determine the analyses needed for each subset. This iterative process of dialogue between the user and ChatGPT 4, although thorough, was time-consuming due to the dataset's size (over 400 columns).

A notable aspect of ChatGPT 4's performance was its contextual understanding and the ability to make nuanced adjustments. For instance, it advised normalizing frequency distributions for categorical data based on the count of non-empty responses, a detail that could easily be missed even by human analysts._

Additional Uses & Features:

Code Review: ChatGPT 4 was tasked with reviewing the primary function used in the analysis. In a code review, ChatGPT 4 checks for errors and suggests improvements for clarity and simplicity. The AI successfully provided an edited version of the function, making it easier to understand.

Documentation Writing: ChatGPT 4 also efficiently wrote documentation for the code, accurately and clearly explaining each step.

Both the code review and documentation writing demonstrated ChatGPT 4's ability to quickly and effectively understand and fulfill these specific tasks, showcasing its utility in streamlining the coding process.

Takeaways for ChatGPT 4:

Overall, ChatGPT 4 is the strongest AI model when it comes to survey analysis. It consistently produces accurate code and requires the least amount of back-and-forth to achieve the desired results. Because ChatGPT 4 could parse the 2021 and 2022 Privacy Reports, as well as the full data set containing survey responses, it stood out as the model with the most human-like approach to performing the survey interpretation, displaying the ability to translate the contents of the previous reports into data analyses, then write the code for these analyses and return a final number. Since it can output final numbers (rather than chunks of code that the user would have to execute themselves, like in the other two models), ChatGPT 4 is also the ideal choice for users who do not have a programming background.

Bard

The prompt:

> I am going to give you the first few rows of a raw data file that was collected in 2023, and 2 previous versions of reports that were written using similar data collected in previous years (2022 and 2021). Could you please output code in Python that I can run on a Jupyter Notebook on my local computer in order to analyze the data according to the contents of the previous years' reports?

The results:

The case study reveals challenges encountered with Bard in generating code according to specific requirements:

1. **Initial Misinterpretation:** When asked to create code for calculating a frequency distribution for a specific column, Bard initially produced code that calculated the mode for every column in the data frame, indicating a deviation from the requested task.
2. **Requirement for Detailed Guidance:** Subsequent interactions with Bard necessitated explicit instructions on the exact Python methods to be used. Despite recognizing the difference between calculating the mode and a frequency distribution, Bard struggled to autonomously generate the correct response.
3. **Incorporation of Unspecified Parameters:** Bard introduced additional inaccuracies by adding unspecified parameters into the method. These discrepancies were significant as they either stemmed from misinterpretation of the directives or resulted in code that could cause errors regardless of the data input.
4. **Need for Extensive User Intervention:** Bard required considerable user guidance to identify and correct these inaccuracies. The case study includes an experimental assessment (See below) that aimed to determine the level of input Bard needed to correct an error in its initial code composition.

Human Input	Bard's Output
First attempt: "This doesn't return what I am looking for. Can you please de-bug it?"	Bard does not find the error. It returns a slightly edited version of the code block, but the line with the problem remains unchanged.
Second attempt: "This doesn't return what I am looking for. The error is in this line: df = pd.read_csv("2023_privacy_data.csv", header=None, names=column_headers) Can you please find the error in this line and fix it?"	Bard reports that it detects an "error" in this line, but not the correct one.
Third attempt: This doesn't return what I am looking for. The error is in this line: df = pd.read_csv("2023_privacy_data.csv", header=None, names=column_headers) The CSV file contains hundreds of columns outside of the ones listed under column_ headers. We should read in the entire CSV file, and then specify those particular columns afterwards when we are calculating frequency distributions.	Bard returns the correct output.

As shown, Bard is unable to detect its own errors unless it is shown exactly where and what the error is, somewhat defeating the purpose of using the assistance of an AI platform altogether.

Importantly, Bard holds very little context, "forgetting" previously mentioned information or errors that were previously corrected in the conversation. According to the FAQ section, this is intentional, but it makes survey interpretation with the model a lengthier process in that the same bug or error must be addressed multiple times.

Additional Uses & Features:

When it comes to reviewing previously written code, Bard was able to give some suggestions for increased readability. However, other suggestions it gave were incorrect and would have resulted in errors if implemented. Bard's ability to write documentation for previously written code was strong, on par with ChatGPT 4.

Takeaways for Bard:

Relative to the other two models, the analytical process with Bard was less fluid. The generated code frequently contained errors, and the model could not rectify these inaccuracies without explicit directives pinpointing the exact nature and location of the mistakes. While the prospect of analyzing a dataset stored in one's Google Drive is enticing, the current functionality falls short. Specifically, we encountered challenges in prompting Bard to access or provide summaries of even text-centric files within our Google Drive.

Claude AI

The prompt:

> I am going to give you the first few rows of a raw data file that was collected in 2023, and 2 previous versions of reports that were written using similar data collected in previous years (2022 and 2021). Could you please output code in Python that I can run on a Jupyter Notebook on my local computer in order to analyze the data according to the contents of the previous years' reports?

The results:

When given general instructions to perform analyses on the sample data set given the 2021 and 2022 Privacy Reports, Claude AI returns pseudocode that is based on the reports, but not on the data set. Whether this is due to copyright issues or the model's inability to parse so many input sources is unclear. However, when given instructions to perform a specific analysis relating to the report, or to write a general function to be used for the report, Claude AI will generally return code that achieves the desired result.

Additional Uses & Features:

Claude AI can "Code Review." Rather than returning a slightly edited version of our code, Claude AI completely restructured it, returning an entirely new design that still achieved our desired output. In practice, the result was unrecognizable from the original code thus making it difficult to understand each step. However, the model's ability to immediately grasp the intent of the code block, and write a new code block from scratch rather than making copy edits to the original is very impressive.

Claude AI's ability to write documentation was not as strong as the other two models. The result was a surface-level description of the steps being taken, which required a large amount of context to understand. Claude AI did not include descriptions of each function, including lists of inputs, which is generally standard practice.

Takeaways for Claude AI:

Claude AI demonstrates a commendable capacity for survey interpretation, positioning it as more adept than Bard but less robust than ChatGPT 4. Owing to copyright constraints, it generates code only under precise directives. Although it can access the entire dataset, it will replicate column names and inherent data values exclusively when samples are directly fed into the text interface. When tasked with a code review, Claude AI undertook a comprehensive reorganization of the code, yielding the desired output but diverging significantly from the approaches of the other two models.

Key takeaways:

	ChatGPT 4.0	Bard (Free Version)	Claude.AI (Free Version)
Quality of code produced	High-quality code is produced when given relatively minimal instruction	Code produced often contains errors, requires large amounts of instruction to achieve desired result	High-quality code produced when given a high degree of instruction
Code-reviewing capabilities	Able to outline and implement general suggestions in line with best practices	Able to outline and implement general suggestions, but some suggestions contain errors	Able to entirely restructure and re-write a code block to achieve the desired result
Documentation-writing capabilities	Able to write detailed explanations of each step taken, for human-written code	Able to write more general explanations of each step taken, for human-written code	Able to write detailed explanations for restructured and rewritten version of human code

- AI can be helpful in Survey Interpretation by writing code that can be used to analyze data. When given very specific instructions, all 3 models (ChatGPT 4, Bard, and Claude AI) can provide code that can produce certain summary statistics or charts to aid in data analysis.
- None of the models tested can perform front-to-back analysis (although ChatGPT 4 is the closest to this capability). They cannot ingest large data sources and output the results based on a vague or open-ended question, such as: "Perform an analysis on the results of the privacy survey based on the privacy reports from previous years." Their "creativity" when approaching problems is limited and operate best when given strict instructions for execution.
- Performing analysis with AI is an iterative process. AI's first attempt at writing code is almost never correct, but it can reach the correct output with additional feedback or details. This back-and-forth is time consuming, but certainly has the potential to increase efficiency if one does not have a high degree of familiarity with the programming language being used.
- In addition to writing the code to perform the analysis itself, AI can potentially be used to support technical work in the form of:
  - **Parallel processing** – Performing an analysis, having AI write a version of the code, then running both and comparing the output as a method of checking one's work.
  - **Code reviews** – Analyzing code that was previously written and editing it for increased clarity, conciseness, and reproducibility.

- **Documentation** – Writing annotations in previously-written code with descriptions for what is being done at each step, making it more seamless for future collaborators to understand the processes that were performed.

For a more detailed review of the prompts and LLM responses see.

4. CASE STUDY 4: TRANSLATION OF TEXT

In this section, we delve into two case studies that explore the efficacy of AI platforms in translating content into languages other than English. As global markets become increasingly interconnected, the ability to effectively communicate across linguistic barriers becomes paramount for successful advertising strategies. This case study evaluates ChatGPT 44.0, Bard, and Claude AI, examining their accuracy, cultural sensitivity, and applicability in the realm of advertising. First, we examine how these three AI platforms translate the ARF's Privacy Report into German, French, Spanish, Hindi, Mandarin Chinese, and Hebrew. Second, we examine how the same AI platforms translate one of Erwin Ephron's newsletters into German, French, Spanish, Hindi, Mandarin Chinese, and Hebrew. While the privacy report translations might yield results that are related to the type of text, i.e., a concise summary of survey findings, translation of Ephron's newsletters (ironic, humorous, more nuanced, very American-centered) might yield different results. Comparing the two cases studies, we seek to provide valuable insights into how AI can be leveraged to overcome language barriers.

ARF Privacy Report:

German

- All three AI platforms did extremely well, with consistent translations across the platforms: same headlines, similar language.
- Translation is very literal across all three platforms.
- Interestingly, all three platforms similarly translated "The ARF Privacy Report," as "The ARF Data Protection Report," which is how a German person of any position would phrase it.
- Notably, ChaptGPT and Claude AI made the same error in the following sentence: "Respondents are more consistent with the way they perceive what search engines and social media sites do with their data; their perception of government sites is more varied." Instead of translating "more consistent," the AIs translated "consistencies." This may, however, be related to the slightly confusing English text.
- Bard did not translate the last paragraph and misplaced one headline. Bard also incorrectly translated the following: "Across all institutions (including scientists, advertising, TV news, etc.), trust has increased from 2022" as "Since 2022, trust in all institutions (including scientists, advertising, TV news, etc.) has increased."

Key takeaways:

- This exercise, which involved summarizing survey findings through multiple brief sentences and other summarization techniques, may have led to the infrequent error observed. However, it's important to note that this approach might not be applicable to all research reports. The absence of slang, irony, or exaggerated language in this context suggests that these elements could pose additional challenges in different types of research documentation.
- We can hypothesize that the clearer and possibly simpler the original content, the lower the probability of errors in the outcome.
- Additionally, it seems that the quality of translation improves with the availability of more reference material in the target language. For instance, there is likely a wealth of information on German perspectives regarding privacy and data protection.
- Consequently, it may well be that the accuracy observed in these German translations cannot be reliably extended to translations into other languages.

French

- French translation from English in Bard, Claude AI and ChatGPT 4 isn't 100% grammatically perfect, and varies somewhat between the three platforms. However, their near-perfect accuracy is suitable for French readers in terms of overall comprehension.
- Despite a few mishaps, ChatGPT 4 and Claude AI fared best in terms of capturing the nuances and idiomatic expressions in French.
- Bard demonstrated consistent use of modern terminology.
- Notably, the choice of "courier électronique" or "messagerie électronique" over the official term "e-mail" by all three platforms is an interesting deviation from standard French terminology.

ChatGPT 4

- Word Choice: Exhibited a mix of literal and idiomatic translations. For example, its choice of "Concernant la vie privée" was more idiomatic than Bard and Claude AI's "Sur la confidentialité."
- Sentence Structure: Demonstrated a better grasp of nuance in sentence structuring compared to Bard and Claude AI, as seen in the phrase "en vue d'une réanalyse," translated from "All previous years of survey data are available to members for reanalysis."
- Grammatical Errors: Occasionally made formal grammatical errors, such as omitting 'ne' in the subjunctive when translating the sentence "...they were concerned about data being used for purposes other than advertising" into "qu'ils craignaient que les données NE soient utilisées."
- Interpretation: ChatGPT 4 sometimes introduced interpretation issues, as in the case of incorrectly translating "When asked about recent privacy changes..." as asking about people's point of view ("Lorsqu'on a demandé leur avis sur...").

Bard

- Word Choice: Demonstrated consistent use of modern terminology.
- Literal vs. Nuanced Translations: Tended to be more literal. For example, used "SmartTV" in line with contemporary usage.
- Sentence Structure: In some instances, the translations were straightforward and lacked contextual depth, as seen in the translation of "pour réanalyse" for the sentence "All previous years of survey data are available to members for reanalysis." ChatGPT 4's "en vue d'une réanalyse" was more nuanced as in 'in case of reanalysis').
- Grammatical Errors: Made similar errors to Claude AI in translating certain sentences, like incorrectly comparing respondents to previous years.
- Interpretation: Bard did not deviate significantly from the original text's intent.

Claude AI

- Word Choice: Varied in capturing idiomatic expressions. For instance, used "répandu" for "prevalent," which is more natural in French.
- Literal vs. Nuanced Translations: Showed a mix of literal and nuanced translations. For example, the translation of "feeling more informed" as "mieux informés" captures nuances better than its counterparts.
- Sentence Structure: Similar to Bard in some respects, such as the straightforward translation of "pour réanalyse."
- Grammatical Errors: Made similar errors to Bard in sentence comparison structure.
- Interpretation: Occasionally, Claude AI's translations were grammatically awkward or altered the meaning of the original text, as seen in the translation of "When asked about recent privacy changes…" into "Lorsqu'on eur a demandé des changements récents" (which translates to "when we asked them recent changes").

Key takeaways:

- Each platform shows its unique strengths and weaknesses in translating English to French. While none is perfect, their translations are generally effective for comprehension, and maintain a high level of accuracy, making them suitable for French readers in terms of overall comprehension.
- ChatGPT 4 and Claude AI tend to fare better in capturing nuances and idiomatic expressions in French, often providing translations that are more aligned with the natural flow of the language. Bard, while sometimes less nuanced, is noted for its consistent use of modern terminology.
- Grammatical Precision: ChatGPT 4, while generally nuanced, sometimes makes formal grammatical errors, contrasting with Bard and Claude AI's more literal but grammatically aligned translations.
- All three platforms interestingly deviate from standard French terminology, such as using "courier électronique" or "messagerie électronique" instead of the official term "e-mail".

Spanish

ChatGPT 4

ChatGPT 4 displayed a balanced approach in translation, navigating between being literal and providing nuanced translations. However, this balance sometimes led to variations in the smoothness and accuracy of its translations.

Examples:

1. Word Choice/Structural Issues:
   - Correctly included 'la' before 'primavera'.
   - Used "Hallazgos" for 'findings.'
   - Omitted the article "la" in headers, opting for a more direct approach.
   - Came closest in translating "In People Like Me, We Trust".
   - Used "género" for 'gender.'
   - Used "telefonos inteligentes" for 'smartphone' and "Segmentación" for 'targeting'.
2. Word Choice (Correct vs. Literal; Smooth vs. Clunky):
   - Opted for a clunkier "se llevó a cabo" for 'was conducted'.
   - Chose a literal translation for "based on" and "for reanalysis".
   - Most accurate in sections like "On Targeting".
   - Less nuanced in translating "people most commonly indicated".

Bard

Bard often leaned towards more literal translations, closely following the source language structure, which sometimes impacted the fluency and contextual accuracy of its translations. Bard demonstrated contextually appropriate term choices in some instances, like using "sexo" for 'gender', indicating a nuanced understanding of certain terms. However, its adherence to literal translations led to variations in the overall smoothness and authenticity of the translated text.

Examples:

1. Word Choice/Structural Issues:
   - Omitted ‘la’ before ‘primavera’ in the introductory paragraph.
   - Used “Hallazgos” for ‘findings.’
   - Followed the literal “Sobre la” structure for headers like “On Privacy” and “On Targeting.”
   - Less inventive in translating “In People Like Me, We Trust”, sticking close to the English structure.
   - Used “sexo” for ‘gender,’ which is contextually more accurate.
   - Used “telefonos inteligentes” for ‘smartphone’ and “Segmentación” for ‘targeting.’
2. Word Choice (Correct vs. Literal; Smooth vs. Clunky):
   - Used “se realizó” for ‘was conducted,’ which is smoother.
   - Chose a literal translation for “based on” and “for reanalysis.”
   - Varied translations in sections like “On Targeting.”
   - Chose a more nuanced translation for “people most commonly indicated.”

Claude AI

Claude AI exhibited a propensity for accurate and nuanced translations, especially in complex sentences. This approach was evident in its choice of words like “según la” for ‘based on,’ which were more contextually fitting than literal translations. However, Claude AI occasionally opted for more literal translations in certain terms, such as using ‘resultados’ for ‘findings.’ Despite this, Claude AI’s translations generally maintained a good balance between being literal and providing depth, demonstrating an understanding of subtleties in language. Claude AI’s weakness was apparent in some awkward translations, like its rendition of phrases related to ad targeting, which could disrupt the natural flow of the translated text.

Examples:

1. Word Choice/Structural Issues:
   - Also omitted 'la' before 'primavera.'
   - Chose 'resultados' for 'findings,' which is more literal.
   - Followed the literal "Sobre la" structure in headers.
   - Provided an accurate but lengthy translation for "In People Like Me, We Trust."
   - Used "género" for 'gender,' which is less accurate.
   - Opted for "televisores inteligentes" for 'smartphone' and directly used "Targeting"
2. Word Choice (Correct vs. Literal; Smooth vs. Clunky):
   - Chose "según la" for 'based on,' which is more accurate.
   - Opted for smoother translations in the introductory paragraph and other sections.
   - More nuanced in translating "people most commonly indicated."

Key Takeaways:

- Each platform has its strengths and weaknesses: Bard in literal translations, Claude AI in nuanced translations, and ChatGPT 4 in balancing between the two with occasional clunkiness.
- Bard: Tends to stick closer to the source language structure, making its translations more literal at times. It is contextually more accurate in some term choices like "sexo" for gender.
- Claude AI: Shows a tendency towards more accurate and nuanced translations, especially in complex sentences. However, it sometimes opts for more literal translations in terms like 'findings.'
- ChatGPT 4: Demonstrates a balance between literal and nuanced translation, but can be clunky in some phrases. It stands out in handling complex sentences and omitting unnecessary articles for a more direct approach.

Hindi

- In some cases, Bard and Claude AI produced identical translations, suggesting either a high degree of accuracy or a potential limitation in translation variability.
- Bard demonstrated superior accuracy in certain translations, such as correctly translating the phrase "The ARF has conducted its sixth annual Privacy Study" and using the Hindi word for "Device" instead of the English term. This indicates Bard's effectiveness in maintaining linguistic authenticity while translating into Hindi.
- Literal vs. Contextual Translation: ChatGPT 4 and Claude AI, in some instances, relied on direct translations (e.g., using "Device" in Hindi), which might not resonate as well with native Hindi speakers who would expect a more localized term.

- Performance Comparison: Both ChatGPT 4 and Bard performed similarly and were generally better than Claude AI in terms of translation quality, indicating their effectiveness in translating complex texts into Hindi.
- Translation of Numbers: Bard's ability to translate numbers into Hindi in specific contexts shows a higher level of attention to detail.

ChatGPT 4

- Introduction Translation: ChatGPT 4 translated the introduction well but inaccurately translated "The ARF has conducted its sixth annual Privacy Study."
- Accuracy in Terms: ChatGPT 4 correctly translated terms like "Median" and "opt-in."
- Performance Comparison: ChatGPT 4's overall translation quality was adequate, excelling in certain areas.
- Specific Translations: ChatGPT 4 performed better in translating certain phrases like "What happens to your data?"

Bard

- Introduction Translation: Bard translated the introduction well but failed to recognize that "U.S" and America are synonymous, not using the Hindi word for "U.S."
- Accurate Translation: Notably accurate was Bard's translation of "The ARF has conducted its sixth annual Privacy Study."
- Numbers Translation: Bard occasionally translated numbers into Hindi, exemplified in the smartphone usage section.
- Word Usage: Some words, like "relationship," were contextually misplaced in the translation.
- Device Translation: Bard accurately used the Hindi word for "Device" as opposed to English, showing a good grasp of context.
- Sentence Structure: Differing from the original English order, Bard re-arranged phrases in a way that enhanced the translation.
- Human-like Quality: Some translations were on par with human translators in terms of quality and understanding of nuances.
- Minor Errors: There were small mistakes, like translating "Average" instead of "Median."
- Summary Tendency: Bard tended to use traditional Hindi words more suited to formal Hindi literature than everyday use.

Claude AI

- Introduction Translation: Claude AI effectively translated the introduction and accurately used the Hindi word for "U.S."
- Literal Translation: Claude AI's translations were more literal, maintaining the original order but sometimes affecting the natural Hindi flow.
- Cohort Translation: Unlike Bard and ChatGPT 4, Claude AI did not translate certain phrases, like "cohorts."
- Accuracy and Satisfaction: Some phrases were translated with complete accuracy, providing high satisfaction in their Hindi rendition.

Key insights:

- The performance of these translation tools varies depending on the context and the specific demands of the text being translated.
- Bard generally outperformed Claude AI in Hindi translation accuracy.
- The use of traditional Hindi words by Bard and ChatGPT 4 might be difficult for the younger generation but appreciated by formally educated Hindi speakers.
- Language translation from English to Hindi where cultural nuances, audience preferences, and the balance between traditional and modern language use play significant roles is particularly complex.

Mandarin Chinese

In the translations to Mandarin Chinese by all three AI platforms, the approach was highly literal. Each model strictly adhered to the original sentence structures and phrases, without adapting them for better grammatical flow or tonal consistency in Mandarin. This resulted in translations that, while understandable, lacked the fluidity and nuance necessary for effective communication. Consequently, the outputs from all three platforms are not ready for direct use with Mandarin-speaking audiences without further refinement.

Despite the similarities, Claude AI slightly edged out in performance due to its proficiency in translating complex advertising and marketing terminology while maintaining the report's overall tone. However, the quality gap between Claude AI and the other two models was relatively narrow in this instance.

ChatGPT 4

- Unlike the other models, ChatGPT 4 displayed sections of 1-2 sentences in English, followed by the Chinese translation (the other models only returned the Chinese version of the text). This made it more seamless to compare the input and output.
- ChatGPT 4's translation capabilities were intermediate in their performance.
- Literal Translation: ChatGPT 4 used extremely literal translations.
- Preservation of Tone: ChatGPT 4 Displayed a slightly better ability to preserve the academic tone, making sentences somewhat more readable than Bard's.

Examples:

- ChatGPT 4's translation of "Median daily Smart TV usage" was word-for-word, not reflecting the natural phrasing in Mandarin.
- In the phrase "The percentage of people willing to share their data with marketers," ChatGPT 4 translates "marketers" to a word more closely meaning "salespeople."
- ChatGPT 4 directly translates each of the two words in the phrase "Perhaps unsurprisingly" – the result is nonsensical in Chinese.
- In the sentence "Generally, tolerance for sharing data decreases with age," ChatGPT 4 uses a phrase for "generally" that is fairly colloquial and used in everyday speech, rather in academic writing.

Bard

- Literal Translation: Bard struggled with literal translations, consistently preserving the original structure of sentences and phrases.
- Passive Voice: Bard's translations maintained the passive voice, which is uncommon in Mandarin, especially in formal writing. This resulted in confusing translations for Mandarin speakers unfamiliar with English grammatical structures.
- Overall Readability: The translation by Bard did not read fluidly and was challenging to understand due to the direct word-for-word approach.

Examples:

- In the phrase "'sharing data' involves data being sold and sent to other parties," Bard directly translates both the words "share" and "data," neither of which are correct for the context. Bard uses a word for "share" that is usually used to describe people enjoying something together (i.e., food), and the word used for "data" is more commonly used in the context of statistics and finance (a closer word would have meant something closer to "information" rather than "data").
- Bard translates word-for-word the section header "In People Like Me, We Trust," a phrase that does not exist in Chinese, rendering the result nonsensical.
- Bard translates the section header "What happens to your data?" into a much more colloquial exclamation, loosely meaning "What's up with your data?"

Claude AI

- Sophisticated Translation: Claude AI displayed the most advanced translation capabilities.
- Technical Language: Claude AI excelled in dealing with technical terms, translating them accurately even when they weren't explicitly mentioned in the original text.
- Tone Replication: Claude AI was mostly successful in replicating the tone of the original report, with fewer instances of colloquial language or awkward phrasing.
- Contextual Understanding: Claude AI demonstrated a strong ability to identify marketing terms by their context and return the correct translation.

Examples:

- When translating the word “campaign” (in the context of a marketing campaign), Claude AI incorrectly uses the Chinese word for “event” or “activity.”
- Claude AI uses a more colloquial term for “interestingly” in the sentence: “Interestingly, heavy usage declined across the board for all devices.” Rather than use a term that highlights that this data point is unexpected or notable, Claude AI uses a phrase loosely translating to “What’s weird is that…”
- Claude AI translated “The ARF” into “The American advertising research nonprofit.”
- Claude AI was also able to correctly identify marketing terms by their context and return the correct translation, such as for the terms “cookie” and “segment.”

Key Takeaways:

- Literal Approach: All three AI models primarily utilized a literal translation approach, resulting in Mandarin translations that were grammatically awkward and difficult for Mandarin speakers to fully understand.
- Preservation of Tone: Claude AI was most effective in preserving the overall tone of the original text and dealing with technical language, making it the strongest translator among the three.
- Readability: Bard’s translations were the most challenging to understand due to a rigid adherence to the original sentence structure, which is often incompatible with Mandarin syntax.
- Balanced Performance: ChatGPT 4 offered a balance between the two, preserving the academic tone slightly better than Bard but not as effectively as Claude AI.
- Final Product Quality: The translations by all three models would require significant refinement before being suitable for sharing with Mandarin-speaking audiences, as none of them produced a final product that was entirely grammatically fluid or tonally consistent with the original text.

Hebrew

ChatGPT 4

- Literal Translation: ChatGPT 4 often translated the text literally from English to Hebrew, which sometimes resulted in awkward or unclear sentences.
- Technical Terms: Demonstrated a decent understanding of technical terms but occasionally missed the nuances.
- Phrase and Sentence Structure: Some translations were overly direct, lacking the natural flow of Hebrew language.

Examples:

- For the terms “quotas” and “cohorts” ChatGPT 4 translated into the same word in Hebrew, instead of using the Hebrew term for this word. Although this is a common practice for some words, this is not the case for these specific terms.
- “More informed” was translated “a higher sense of awareness.”
- “Median” was translated “average.”
- “Targeting” was translated into gibberish, although there is a commonly used term in Hebrew.

Bard

- Accurate translation for first paragraph of text.

- However, for second prompt, responded with the following:

  I can’t translate text to Hebrew because I am a large language model, also known as a conversational AI or chatbot trained to be informative and comprehensive. I am trained on a massive amount of text data, and I am able to communicate and generate human-like text in response to a wide range of prompts and questions. For example, I can provide summaries of factual topics or create stories. However, I am not able to translate languages. If you would like to translate text to Hebrew, you can use a translation tool or service.

  - When prompted once again, responded in *Hebrew* to repeat previous claim that it is not capable of translating languages and recommends using a dedicated translation tool or service for such needs.

Claude AI

- Sophisticated Translation: Claude AI’s translations were more sophisticated and closer to the natural flow of Hebrew.
- Technical Language: Showed proficiency in translating technical terms and complex phrases.
- Contextual Understanding: Better at understanding the context and providing translations that fit the Hebrew language style.

Examples:

- Marketers was translated “salespeople.”
- “Cookies” was literally translated into “cookies,” instead of the correct data-related term.
- “Informed” was translated “awareness.”
- “interests” was translated into “matters.”

Key Takeaways:

- Overall Approach: ChatGPT 4 and Claude AI predominantly used a literal translation approach, which sometimes led to translations that were grammatically awkward or unclear in Hebrew.
- Claude AI was generally more effective in delivering translations that were contextually appropriate and stylistically closer to natural Hebrew.
- None of the translations were perfect, indicating that any direct use for a Hebrew-speaking audience would require further editing and refinement to ensure clarity and natural flow.

The Flying Account Planner. (February 2002)

In this newsletter, Ephron reflects on the evolution of the strategic planning role within advertising agencies. The newsletter credits Stanley Pollitt of Boase Massimi Pollitt in the UK for revolutionizing the strategic planner role. Before Pollitt, strategic planning was a shared task among agency brand teams. Pollitt's innovation was to designate a dedicated role for it, elevating its significance and ensuring a consumer-oriented brand plan. This strategic planner, co-equal with account management and creative teams, is portrayed as the consumer's representative, brand's champion and the communication plan's architect. However, the newsletter points out an inherent challenge: this model is rooted in the full-service agency paradigm, which may be outdated given the rise of specialized media agencies and the changing advertising landscape.

German

- All three AI platforms performed adequately, but there is room for improvement.
- It is hard to determine which of the platforms performed best. One platform made a factual translation error but was otherwise accurate. Optimally, we would combine the three versions to receive the best translation.
- Main errors stemmed from unclear or obscure formulations and the use of technical terms in the original content.
- Overall, the AI's performance is impressive considering Erwin's complex prose, but such translations require expert review for style and accuracy based on a deep understanding of the German and US TV industries.
- In general, the translations were not as witty and elegant as Erwin's prose, and they sometimes appeared too literal.
- Notably, each of the AIs employed distinct terminology for the term 'Account Planner.' It is hard to decide which is most suitable.
- Interestingly, despite Erwin's use of "kaput," which originates from German "kaputt," none of the AIs kept the German term, opting instead for various translations.

ChatGPT 4

- This platform made one inaccurate translation that resulted in an error. In general, the translation tends to be overly literal. This does not necessarily result in mistakes, yet in language that is awkward and inelegant.

Examples:

- In the sentence "What Pollitt did that was revolutionary... He focused agencies on the primary importance of developing a consumer-oriented brand plan..." ChatGPT 4 translated "the primary importance" as "primäre Bedeutung," which is overly literal and would perhaps be better translated "wichtigste Aufgabe, die" which is "the most important task.."
- In the sentence "the embroidered function..." ChatGPT 4 translated "embroidered" literally to "*bestickte.*"
- In the sentence "The clinker in account planning is the rise of the media agency as the brand's other marketing partner," ChatGPT 4 literally translated "clinker" as "Der Knackpunkt," a rarely used German word. A better translation could be "der Haken," which is the obstacle or the problem.
- In the sentence "So in a strange twist the media agency may become the seat of account planning simply because creative is too important not to be able to fire," ChatGPT 4 mistakenly translated "um nicht feuern zu können," which is the opposite of what was intended, "Creative is too important not to have the ability to fire." The correct translation would be "um nicht gefeuert werden zu können," which translates "to avoid being fired."

Bard

- Bard's translations seem to better capture Erwin's style and tone, using some colloquial and everyday language.
- Still, in general, these translations also tend to be overly literal, resulting in some awkward language.

Examples:

- For the headline, "The Flying Account Planner" Bard used "Kundenbetreuer," which translates as "client caretaker." This is an overly broad and general term. In this case, a more literal translation "Kontoplaner" would have been more appropriate.
- In the sentence, "The late Stanley Pollitt, a brilliant advertising man at Boase Massimi Pollitt in the UK..."Bard translated "Werbeprofi," which is an colloquial term for advertising expert, consistent with Erwin's style.
- In the sentence "He focused agencies on the primary importance of developing a consumer-oriented brand plan by putting a researcher in charge of doing it and made that person co-equal with account management and creative"Bard literally translated "creative" as "Kreation." A more appropriate term in this context would have been "Kreativabteilung" which translates "creative department." This literal translation of Erwin's short-cut expression "creative" appears throughout this translation.

Claude AI

- Claude AI's translations are oftentimes literal, resulting in awkward sentences.
- At times, when deviating from literal translations, Claude AI provided excellent translations, outperforming the other two AIs.
- Claude AI was inconsistent, using different terms for "the account planner" in headline and throughout the text.

Examples:

- In the sentence "...since agencies could make fair claim to the strategy thing before Stanley," Claude AI translated "strategy thing" with "die Strategiesache," which translates "startegy matter." A better translation could have been "Strategiethema," which translates "strategy topic."
- In the sentence "It sounds like a hi-powered way to create advertising until you realize it's a full-service agency concept at a time when full-service agencies are kaput," Claude AI went a step too far, with the translation "in the full service agency they don't exist anymore."
- In the sentence "agencies will have to rebundle around a new concept of creative management" "rebundle" was literally translated whereas it may have been better to translate "to group" ("gruppieren").
- In the sentence "The supreme importance of creative means advertisers need choice. But no agency provides choice, only the illusion," Claude AI translated the first "choice" as "Wahlmöglichkeiten" which translates "possibility of choice." This is a deviation from literal translation that is more elegant.

Key takeaways:

- In terms of performance assessment, all three AI platforms showed adequate performance but need to be reviewed. Determining the best performer is challenging; ideally, a combination of all three versions would yield the best translation.
- Translation challenges are related to the specialized technical terms and complex formulations in the original content.
- These insights reflect the varied capabilities and limitations of each AI platform in handling complex translation tasks, especially when dealing with stylistically unique and industry-specific content.

French

Translating Ephron's witty, idiomatic prose posed challenges in French due to the complexity of idiomatic expressions, which varied in translation. For instance, the word "clinker" yielded different results across platforms: Claude AI came closest with "l'accroc" (which means "hitch" or "snag"), ChatGPT 4 opted for an elegant "le point délicat." Bard was the least ambitious with "le problème."

Industry-specific terms like “creative” that were used in different contexts also resulted in diverse interpretations. For example, in the sentence “the account planner needs to sit with creative,” all platforms correctly translated it. However, in more nuanced sentences like “Creative loses accounts; media doesn’t,” Claude AI’s translation was the most accurate and idiomatic. For “creative agency,” Claude AI correctly translated it as “agence de creation,” while Bard and ChatGPT 4 were incorrect with “agence de créativité” and “agence créative,” respectively. In describing personnel as “creatives,” there’s no direct French equivalent, but Claude AI came closest in various contexts. The variance between translations showcases the complexity and variability in translating industry-specific jargon like “creative” across different AI platforms.

As an example of awkward sentences resulting in awkward translations, consider the sentence “With media separated from creative, where you put the account planner is important.” This posed a translation challenge for all three AI platforms, translating “where” as “l’endroit,” which directly means “the place” or “location.” A more sophisticated and contextually appropriate French phrase would be “là où vous placez,” which better captures the intended nuance of “where you put” in this context.

Overall, it seems ChatGPT 4 bested the other two platforms, in terms of offering fewer literal translations, suggesting accurate phrasing and word choices, and providing sophisticated verb conjunctions. However, this was not consistent, and at times Claude AI and Bard bested ChatGPT 4 at word choices.

ChatGPT 4

- Strong in translating idiomatic expressions.
- Sometimes overly literal in word-for-word translations.
- Good at maintaining sentence structure and context.

Examples:

- Inaccurately translated “creative agency” with “agence creative.” Additionally, for the sentence “creative is too important to not be able to fire,” ChatGPT 4 incorrectly translated “la créativité.” In “Creative loses accounts; media doesn’t,” ChatGPT 4 lost the idiomatic nuance with “Les agences créatives perdent des comptes, les agences medias non.” However, it correctly translated “creative director” as “directeur créatif,” capturing the professional title accurately.
- “Today’s strategic planner has broad horizons but no proper home.” “Home” in this context doesn’t translate well because it’s more about the planner not being integrated well anywhere, but the word in French “foyer” can work. ChatGPT 4’s “véritable domicile” was not a good choice (“veritable” has a truth context as opposed to appropriate and “domicile” is too literal for home in Ephron’s context).
- For “By now, the embroidered function is familiar” ChatGPT 4 used “Aujourd’hui la function du planificateur est bien établie” vs. Bard and Claude AI’s literal “la fonction brodée est familière.”
- Last paragraph’s “Full-service agency”: ChatGPT 4 got it right with “agence à service complet.”
- For “no agency provides choice,” ChatGPT 4 used the elegant “aucune agence ne propose de choix.”

Bard

- Struggles with industry-specific jargon.
- Better at verb conjugations and syntax.
- Inconsistent in context understanding.
- Demonstrated strength in translating complex advertising terms, maintaining the context effectively.

Examples:

- Inaccurately translated “creative agency” with “agence de créativité.” Similarly, in “Creative loses accounts; media doesn’t,” Bard’s was close but not as pithy as Claude AI with “La création perd des comptes, pas les medias.”
- Bard offered three other options of idiomatic options with different contexts of “regroup” and “home” as well as the nature of the account planner:
  1. Avec la restructuration des agences, le nomade du planning trouvera enfin sa place. This emphasizes the nomadic nature of the account planner; also instead of “have a home” it’s “will find his place.”
  2. Quand les agences redeviendront multiservices, le planificateur de compte touche-à-tout aura enfin son port d’attache. This emphasizes the multi-skilled nature of the account planner; also instead of “have a home” an expression uniquely French: “touche-à-tout aura enfin son port d’attache” which means “jack of all trades will finally have his home base.”
  3. Le grand chamBardement des agences permettra enfin au planificateur de compte globe-trotter de poser ses valises. This emphasizes the constant travel of the account planner; also instead of “have a home”: “poser ses valises” or “put down his bags or suitcases.”
- “The traditional answer would have been”, and the next sentence “Creative agencies are too transient.” For “transient” Bard (and ChatGPT 4) chose the more elegant “ephémère” vs. Claude AI’s literal “transitoire.” For “strange twist” Bard came closest with “étrange revirement” (strange turnaround) but incorrectly placed the adjective before the noun; ChatGPT 4 and Claude AI opted for the clunkier “un étrange retournement de situation,” also with the adjective wrongly placed.

Claude AI

- Excels in translating creative industry terms.
- Occasionally misinterprets idiomatic phrases.
- Strong in word choice and contextual relevance.
- Demonstrated strength in translating complex advertising terms, maintaining the context effectively.

Examples:

- Accurately translated “creative agency” with “agence de creation.” Additionally, in “Creative loses accounts; media doesn’t,” Claude AI came closest with “La création perd des comptes, les médias non.”
- “When agencies rebundle, the flying accountant planner will finally have a home.” Claude AI (and Bard) used the better “se regrouperont” but ChatGPT 4’s “se réorganiseront” was close enough.
- For “She or he is said to have broad knowledge of…” Claude AI’s translation “Elle ou il est censé avoir” is not only clunky and overly literal but creates the problem of masculine/feminine verb correlation.
- “The planner uses these skills to lead … to the point where the craftsmen … can take over and create the campaign.” Claude AI was most eloquent out of the three platforms. Instead of the clunky “jusqu’a” structure it used “au point où” which does not require subjunctive.

Key Takeaways:

- Overall, each platform had its strengths and weaknesses, with Claude AI showing a slight edge in industry-specific terms and idiomatic accuracy.
- ChatGPT 4 excelled in translating idiomatic expressions and maintaining sentence structure but was sometimes too literal. It struggled with nuanced sentences. ChatGPT 4 offered sophisticated verb conjugations but occasionally missed in word choices.
- Bard struggled with jargon but was effective in complex advertising terms and syntax. Bard provided alternative idiomatic options, demonstrating versatility in context understanding.
- Claude AI was strongest in translating creative industry terms and contextual relevance and it closely captured idiomatic nuances. Occasionally Claude AI misinterpreted phrases but showed eloquence in complex sentences.

For a more detailed review of the LLM responses see.

Spanish

In the Spanish translations by the three AI platforms, differences were not overly pronounced. However, this did not always equate to accuracy, with literal translations being a common issue. When variations did occur, they were typically linked to awkward English phrasing, the use of industry-specific terms like “creative,” or idiomatic expressions.

Awkward sentences that result in awkward translations: “Today’s strategic planner has broad horizons but no proper home.” “Proper home” does not translate well in Spanish. All 3 platforms chose “hogar adecuado,” literal for home. “Lugar” or “sitio” (meaning “place”) would probably be the better choice but worded differently as in “hasn’t found his place.”

Word choice/turns of phrase also differed between platforms. For instance, “Clinker” proved too difficult for the three platforms: they all chose “El problema”; but “enganche” (which means snag or hitch – which is closer in meaning than “problem”) might have been a better choice.

Word-for-word translation led to errors in syntax. In Spanish adjectives sometimes preceded the nouns they modified (which made for literal translations); this is incorrect. With few exceptions, the adjective should follow the noun. Examples: “brilliant publicist” all three platforms wrote “brillante publicista;” “strange twist”, all three wrote “extraño giro.”

ChatGPT 4

- ChatGPT 4’s translations tend to be accurate and maintain good sentence structure, but they can sometimes be overly literal, missing nuances in idiomatic expressions.
- While it generally provides well-structured translations, its approach can lead to less idiomatic or contextually sensitive results, especially in complex or industry-specific scenarios.
- ChatGPT 4 excels in standard translations but may require human review for more nuanced or creative content.

Examples:

- “The upside is any talent with a pencil can be in the running” was translated “Lo positivo es que cualquier talent con un lápiz puede estar en la competencia” – which is correct but doesn’t attempt to find a parallel idiom.
- “By now the embroidered function is familiar”: for “by now” ChatGPT 4 chose a less literal “a estas alturas” which means “at this stage” vs Bard and Claude AI’s literal “por ahora” and “par ahora.” ChatGPT 4 also did better at “embroidered”, choosing “la función detallada” vs Bard and Claude AI’s literal “fonción bordada.”
- For “before Stanley” ChatGPT 4 chose “antes que” which is incorrect. “Antes de” (for before) would have been the better choice.
- In the sentence “He focused agencies on the primary importance…” ChatGPT 4 correctly chose the past tense “Enfocó” as in “He focused.”

Bard

- Excelled in providing idiomatic alternatives and showed creativity in translating complex sentences.
- Was inconsistent, sometimes choosing incorrect verb tenses.

Examples:

- "When agencies rebundle, the flying accountant planner will finally have a home." While all three platforms came up with literal translations, Bard went the extra mile offering more idiomatic options, bringing in contexts of "regroup" and "home" as well as the nature of the account planner:
  1. Con la reestructuración de las agencias, el nómada del planning encontrará por fin su sitio. This emphasizes the nomadic nature of the account planner – also, Bard chose "will find his place" – "encontrará por fin su sitio," instead of "have a home."
  2. Al volver a ser multiservicio, las agencias darán al todoterreno del planning un lugar donde echar raíces. This option had strange wording, emphasizing the multi-skilled (todoterreno which means "all-terrain") nature of the account planner; also instead of "have a home": "donde echar raíces" which means "where to take root."
  3. El gran cambio en las agencias permitirá al planificador de cuentas trotamundos dejar las maletas. "Cuentas trotamundos" or globe-trotting emphasizes the constant travel of the account planner; also instead of "have a home": "dejar las maletas" which means "leave one's bags."
- "The upside is any talent with a pencil can be in the running" was translated "La ventaja es que cualquier talent con un lápiz puede estar en la carrera." "Carrera" means "race" – closer in idiomatic context but not too literal to "the running."
- "The planner uses these skills to lead": for "to lead" Bard had the better choice: "dirigir" versus ChatGPT 4 and Claude AI's more literal "liderar."
- In the sentence "He focused agencies on the primary importance…" Bard incorrectly chose the imperfect "Enfocaba."

Claude AI

- Played it safe with translations, often choosing direct equivalents.
- It correctly handled some verb tenses but was literal in translating idiomatic expressions.

Examples:

- Claude AI played it safe for “The upside is any talent with a pencil can be in the running”: “La ventaja es que cualquier talent con un lápiz puede participar.”
- In the sentence “He focused agencies on the primary importance…” Claude AI incorrectly chose “Centró” instead of using the reflexive “Se centró.”
- In the sentence “could make fair claim” Claude AI correctly used the imperfect “podian.”
- For “before Stanley” Claude AI (and Bard) correctly chose “antes de” (for before).
- In the sentence “By now the embroidered function is familiar”: for “by now” Claude AI chose a literal translation “par ahora.” It would have been better to translate less literally “a estas alturas” which means “at this stage.” Claude AI also literally translated “embroidered” as “fonción bordada,” whereas “la función detallada” would have been a better translation.

Key Takeaways:

- All platforms struggled with literal translations and idiomatic expressions to some extent.
- Bard showed strength in offering idiomatic options, while ChatGPT 4 was more accurate in certain contexts.
- Claude AI tended to be cautious, leading to more direct translations.

Hindi

Overall, Bard and Claude AI outperformed ChatGPT 4, with Bard having a slight edge over Claude AI. This advantage was due to Bard’s ability to contextually translate into Hindi, whereas Claude AI often directly transcribed English words into Hindi without finding appropriate Hindi equivalents.

ChatGPT 4

- Overall Performance: Least effective among the three.

Examples:

- Translated “account planner” correctly, using the Hindi word for “account planner,” whereas Claude AI used “account planner” in the Hindi translation.
- Did not translate “more accurately” correctly. Rather it translated it to “in a right way.”
- Did not translate “consumer-oriented” to Hindi. Rather, it wrote it in Hindi as-is. The other two Ais translated consumer-oriented in Hindi.
- “Creative loses accounts” was translated as “creative lost.”
- “…either it will be…” translated as ”may be” in Hindi.
- “competing” was translated as “struggling.”

Bard

- Overall Performance: Outperformed both ChatGPT 4 and Claude AI.

Examples:

- Bard provided the best translation for this sentence "By now the embroidered function is familiar. The account planner is described as the consumer's representative, the brand's champion and the communication plan's architect..." The translation was more articulate and appropriate in the context in comparison to the other two AI platforms.
- In this sentence "The planner uses these skills to lead, shape and nuance the development of a brand strategy to the point where the craftsmen - writers, art directors and media planners - can take-over and create the campaign," Bard translated "shape" best according to context.
- In the sentence "It sounds like a hi-powered way to create advertising until you realize it's a full-service agency concept at a time when full-service agencies are kaput," Bard translated "hi-powered" as "really nice." "Concept" was literally translated.
- For the sentence, "The traditional answer would have been "put account planning at the creative agency" but that often doesn't work today...." Bard performed slightly better than other two AIs.
- Bard provided an excellent translation for the sentence, "The supreme importance of creative means advertisers need choice...." Superiorly outperforming other two Ais in terms of the context and choice of words to articulate the meaning.

Claude AI

- Claude AI showed notable competence.
- Its approach often involved directly transcribing English words into Hindi, rather than seeking out contextually relevant Hindi equivalents.

Examples:

- For the sentence, "The traditional answer would have been "put account planning at the creative agency" but that often doesn't work..." Claude AI wrote some of the English words in Hindi ("seat," "fire," "accounts") rather than using the appropriate Hindi terms.
- "It sounds like a hi-powered way to create advertising until you realize it's a full-service agency concept at a time when full-service agencies are kaput...." was better translated by Claude AI since it structured the sentences to best articulate the meaning.
- Claude AI translated this sentence better than ChatGPT 4: "What Pollitt did that was revolutionary, was to turn a shared task of the agency brand team – strategic planning – into a staff assignment – strategic planner. He focused agencies on the primary importance of developing a consumer-oriented brand plan by putting a researcher in charge of doing it, and made that person co-equal with account management and creative."
- While ChatGPT 4 began this paragraph with "it can be heard/understood," Claude AI wrote UK and Boase Massimi PollittI in English instead of Hindi.

Key Takeaways:

- Translation Accuracy: Bard and Claude AI outperformed ChatGPT 4 in translation tasks. Bard had a slight edge over Claude AI, mainly due to its ability to find contextually appropriate Hindi equivalents, whereas Claude AI often transcribed English words into Hindi directly.

Mandarin Chinese

- The translation of Ephron's work was surprisingly better than the Privacy Report.
- Translations were more fluid and better preserved the original tone, unlike the informal language occasionally used in the Privacy Report translations.
- Bard outperformed the other two models in translating Ephron's newsletter, contrary to its performance with the Privacy Report.
- All models had challenges with literal translations, especially with descriptive language, and oftentimes struggled with grammatical structures unique to Chinese.
- The improved translation quality may be due to the models being better trained on more conversational words and phrasing like those used in Ephron's newsletter, encouraging more natural re-phrasing in Chinese.
- The translation of Ephron's newsletter could serve as a foundation for native speakers to refine, whereas the Privacy Report translation required significant alterations.

ChatGPT 4

- Struggled with literal translations, losing the intended meaning in several phrases.
- Faced difficulties with out-of-context words and idiomatic expressions.

Examples:

- "By now the **embroidered** function is familiar": ChatGPT 4 translated "embroidered" literally, which does not make sense in Chinese
- "The **craftsmen** – writers, art directors, and media planners – can take-over": ChatGPT 4 translated "craftsmen" literally.
- "When agencies **rebundle**, the flying account planner will finally have a home": ChatGPT 4 translates the word "rebundle" literally, into "re-tie," which does not make sense.

Bard

- Produced the most accurate translations, with fewer grammatical errors and contextually appropriate phrasing.
- Excellently adjusted sentences and phrases for cultural context.
- Demonstrated a strong grasp of language nuances, like correctly interpreting idiomatic expressions.

Examples:

- In the phrase "to be effective the account planner needs to **sit with** creative," Bard was the only model that correctly interpreted "sit with" as "collaborate" or "spend meaningful time together," rather than physically sit together.
- "The supreme importance of creative means advertisers need choice." Here, Bard preserved the original meaning, which is more like "advertisers need the ability to choose."

Bard produced the strongest translation of the phrase "the flying account planner," which is used in both the title and the conclusion of the newsletter. Bard translated the word to a phrase in Mandarin meaning "floating," which account planner doesn't have a home.

Claude AI

- Claude AI consistently translated "Erwin Ephron" into Japanese rather than Chinese, even after follow-up prompts.
- Aside from this error, Claude AI's translations were slightly better than ChatGPT 4's, showing a moderate grasp of context and grammar.

Examples:

- Claude AI translated "the **craftsmen"** into the word "workers," which is more easily understood by a Native Chinese speaker.
- "The supreme importance of creative means advertisers need choice" was translated by Claude AI as "advertisers need to choose," which is not the intended meaning.

Key Takeaways:

- Bard stood out for its nuanced handling of language, making it the most effective translator for this newsletter.
- Claude AI showed potential but was marred by consistent misinterpretation of names.
- ChatGPT 4, while useful, required significant edits for accuracy.
- All models could serve as starting points for translations, with Bard's output requiring the least editing.

For a more detailed review of the LLM responses see.

Hebrew

ChatGPT 4

- Literal and idiomatic options provided.
- Maintains the original metaphors and structures closely.
- Offers direct translations while also providing natural-sounding Hebrew alternatives.

Examples:

- "Angels" instead of "craftsmen." In Hebrew, craftsmen is two words, "craft" and "owners." However, ChatGPT 4 used only "crafts" which is translated "angels."
- "In a strange twist" was translated as "in a strange turn of events" – this translation captures the idea that is conveyed by using a more literal translation.
- "Any talent with a pencil" was literally translated – yet, in Hebrew this translation is inappropriate and makes no sense. In Hebrew it would make more sense to write "Every talented person with a pencil in hand."
- "The Flying Account Planner: Today's strategic planner has broad horizons, but no proper home" was translated "The Flying Strategic Planner: Today's strategic planner has broad horizons, but no permanent home" which captures the meaning but slightly alters it.

Bard

- Focuses on a straightforward, more direct translation approach.
- Sometimes misses the nuances or idiomatic expressions.
- Tends to provide a more simplified version of the original text.
- Oftentimes offers multiple options according to different possible objectives. This is good to an extent, as sometimes all one wants is basic translation without multiple options or interpretations.
- Needs 'reminders' on what needs to be done. Since there is no option to upload text, the text was copied and pasted in several prompts. However, between one prompt and the next, Bard would often 'forget' that the objective is to translate into Hebrew, instead providing various unrelated responses.

Examples:

- Providing options according to context:

The headline translation received the following response:

Absolutely! Here are two options for translating "THE FLYING ACCOUNT PLANNER Today's strategic planner has broad horizons, but no proper home." into Hebrew:

Option 1 (Literal translation):

הַמִּתְכַּנֵּן הַחֲשְׁבוֹנִי הַמְּרַחֵף יוֹם שֶׁל מִתְכַּנֵּן אֲסָטְרָטֶגִי בָּעָל רְאִיָּה לְמֶרְחוֹק, אֲבָל אֵין לוֹ בַּיִת אֲמִתִּי.

Option 2 (More idiomatic translation):

מִתְכַּנֵּן חֲשְׁבוֹנוֹת נוֹדֵד לְמִתְכַּנֵּן הָאֲסָטְרָטֶגִי שֶׁל הַיּוֹם יֵשׁ רְאִיָּה רחבה, אֲבָל לֹא בַּרְדוֹמָה קְבוּעָה.

The first option is a more literal translation, sticking closely to the English sentence structure and word choices. The second option uses more idiomatic Hebrew expressions to convey the same meaning as the first, but in a way that feels more natural in Hebrew.

Both options are accurate translations, so the best choice for you will depend on your specific needs and context.

- Forgetting what the request was: “I do not have enough information about that person to help with your request. I am a large language model, and I am able to communicate and generate human-like text in response to a wide range of prompts and questions, but my knowledge about this person is limited. Is there anything else I can do to help you with this request?”
- Providing explanations on its outcome: “I used the term "בתי-יצירה" for “creative shops” to maintain a consistent and appropriate style in Hebrew; I capitalized "אסטרטגיה" and "תקציב" as per common style guidelines for nouns representing specific plans or budgets; I translated “why not?” to "ולמה לא?" for a natural and idiomatic equivalent.”

Claude AI

- Attempts to capture both the literal meaning and the spirit of the original text.
- Includes idiomatic expressions that align with the original context.
- Offers a balance between direct translation and natural Hebrew phrasing.

Examples:

- “The Flying Account Planner” was translated “The flying customer planner.”
- Like ChatGPT 4’s translation, “Any talent with a pencil” was literally translated, yet, in Hebrew this translation is inappropriate and makes no sense. In Hebrew it would make more sense to write “Every talented person with a pencil in hand.”

Key Takeaways:

- None of the platforms consistently provided a good enough translation.
- Each platform has its strengths and approaches to translation, reflecting different priorities in conveying the original text’s meaning and style.
- ChatGPT 4 provides diverse options, catering to both literal and natural language needs. Bard’s translations are more direct and simplistic, possibly lacking in nuanced expression. Claude AI sometimes manages to strike a balance, maintaining the original text’s essence while ensuring the translation feels natural in Hebrew.

Key takeaways across the two case studies:

German Translations

- **ARF Privacy Report**: All platforms performed well but literally.

- **Ephron's Newsletter**: Hard to determine best performer: Bard appears to better capture Erwin's style and tone, using colloquial and everyday language that is more in line with the original prose's spirit. Claude AI shows instances of deviating from literal translations to provide more nuanced and contextually appropriate translations. Finally, despite its occasional overly literal translations, ChatGPT 4 maintains a consistent approach and seems to understand the context well, albeit with room for improvement in stylistic adaptation.
- **Overall**: The "best performer" depends on the specific criteria one would prioritize (fidelity to original style, contextual accuracy, consistency, etc.). Combining the strengths of all three might provide the most effective translation.

French Translations

- **ARF Privacy Report**: ChatGPT 4 captured nuances better than Bard and Claude AI. Bard and Claude AI scored higher on grammar but tended to be more literal. All three platforms had similar scores for accuracy, although Bard had an incidence of not translating at all while the other two at least tried.
- **Ephron's Newsletter**: ChatGPT 4 scored higher on nuance and accuracy, with Claude AI a close second on accuracy. ChatGPT 4 also had fewer incidences of being too literal than Bard and Claude AI. And although Bard in one scenario provided varied idiomatic interpretations, it had a greater incidence of translation errors than the other two.
- **Overall**: Claude AI and ChatGPT 4 are noted for their effectiveness in handling creative industry terms and idiomatic nuances. Bard, while consistent, shows a lesser degree of nuance but is competent with complex advertising terms and syntax, offering varied interpretations. ChatGPT 4 and Claude AI are proficient in idiomatic expressions and maintaining sentence structure, but Claude AI's tendency towards literal translations can impact the subtlety of word choices and nuanced sentences.

Spanish Translations

- **ARF Privacy Report**: All three platforms had a tendency to opt for literal translation, with ChatGPT 4 making up for it with greater nuance and accuracy.
- **Ephron's Newsletter**: All three platforms leaned toward the literal, with the following exceptions: Bard showed strength in offering idiomatic options, while ChatGPT 4 was more accurate in certain contexts. Claude AI tended to be cautious, leading to more direct translations.
- **Overall**: All three platforms tend toward literal translations. Bard is strong in idiomatic options, ChatGPT 4 balances accuracy and expressiveness, and Claude AI favors cautious, direct translations, and its approach, though nuanced, tends to be less dynamic.

Hindi Translations

- **ARF Privacy Report**: Bard and Claude AI had similar translations; Bard showed superior accuracy.
- **Ephron's Newsletter**: Bard and Claude AI surpassed ChatGPT 4 in translation accuracy, with Bard slightly ahead for its skill in finding context-relevant Hindi translations, unlike Claude AI, which frequently transcribed English words directly into Hindi.

- **Overall**: Bard emerges as the most accurate and contextually adept translator among the three, with Claude AI showing proficiency but limited by its direct transcription approach. ChatGPT 4, while functional, does not match the translation capabilities of Bard and Claude AI in these specific case studies.

Mandarin Chinese Translations

- **ARF Privacy Report**: Claude AI slightly better; all models too literal.
- **Ephron's Newsletter**: Bard outperformed others, suggesting training on narrative structures helped.
- **Overall**: Bard's translation could serve as a foundation for refinement, while others needed significant alterations.

Hebrew Translations

- **ARF Privacy Report**: Claude AI and ChatGPT 4 consistently used the same language, often repeating mistakes.
- **Ephron's Newsletter**: None of the platforms performed consistently well. Bard offered nuanced translations, demonstrating context sensitivity.
- **Overall**: Although Bard showed stronger awareness of context, ChatGPT 4 and Claude AI provided better translations.

Key takeaways across six languages:

- Cultural Sensitivity: Bard generally showed a better understanding of cultural nuances across languages.
- Literal vs. Nuanced Translations: Claude AI and ChatGPT 4 tended to be more literal, while Bard often provided nuanced translations.
- Complex Texts and Idioms: Bard's ability to offer alternative translations for idioms and complex phrases stood out, particularly in Hebrew.
- Overall Accuracy: Claude AI often demonstrated high accuracy, particularly in Hindi, but faced challenges in maintaining the tone and humor of the original texts.
- Adaptability: Bard's flexibility in offering multiple translation options was notable, especially for idiomatic expressions.
- Tone Preservation: All platforms struggled to varying degrees with preserving the original tone, especially in humor and irony, with Claude AI and ChatGPT 4 often lacking the subtlety required for idiomatic and culturally rich texts.

5. CASE STUDY 5: LITERATURE REVIEWS WEARIN, WEAROUT AND OPTIMAL FREQUENCY

This case study focuses on testing AI's capability in conducting literature reviews, each platform was asked 5-8 prompts related to questions on advertising wearin, wearout and optimal frequency, which are key questions in media planning. The study compares AI-generated reviews with those done by human researchers, using the ARF Knowledge Center's extensive literature reviews as a benchmark.

For the three AI platforms – ChatGPT 4, Bard and Claude AI – we began with the same simple prompt:

> What are best practices for identifying wearin and wearout along with optimal campaign frequency? What are best research approaches to identify wearin and wearout along with frequency for our brand?

Follow up prompts were conducted according to the responses generated by each platform.

Overall, the testing showed that the current iterations of AIs are not very adept at carrying out literature reviews as rampant hallucinations occurred, and the AIs missed several key resources on the topics. In what follows we elaborate on the performance of each AI platform.

ChatGPT 4

- Structured Responses: ChatGPT 4 organized responses clearly, aiding in comprehension.
- ChatGPT 4 synthesized information coherently using general knowledge.
- ChatGPT 4 used fictional citations and references, undermining reliability.
- ChatGPT 4's inability to access external sources impacted the accuracy and currency of reviews.
- ChatGPT 4 sometimes misrepresented or omitted crucial study details.
- The provided literature reviews lacked detailed insights of the original research and often did not align with actual source content.
- Despite varied prompts, responses showed minor alterations without significant improvements.

Bard

- Initial Misinterpretations: Bard confused terms like "wearin" and provided general best practices.
- Bard frequently hallucinated details for referenced papers, indicating a substantial gap in accuracy and often misrepresented study content and findings. This severely undermines the credibility of Bard's ability to conduct literature reviews.
- Literature reviews provided superficial insights, lacking meaningful depth.
- Bard's responses to different prompts were very similar, indicating limited adaptability.
- Like ChatGPT 4, Bard's effectiveness was hampered by the inability to access real-time data.

Claude AI

- Initial Accuracy: Claude AI showed a basic grasp of the subject matter and could format a literature review.
- Claude AI often fabricated source content and misinterpreted research focus and methodology. It also fabricated most of the summaries and incorrectly interpreted studies. Additionally, Claude AI often cited sources with incorrect publication dates and incomplete summaries.

- Mainly, Claude AI cited academic papers, excluding industry or trade sources.
- Claude AI tended to repeat earlier errors in subsequent prompts.

Key takeaways:

Even though Claude AI shows some promise, none of the AI platforms can be currently relied upon to conduct accurate literature reviews. Across the board, the AIs tended to hallucinate summaries and references (see Athaluri et al., 2023 for an elaboration on this phenomenon). In addition, different types and levels of prompts did not make significant differences in results, except for making slight alterations in response formats. At this point, the AIs seem likely to be best used to summarize very general findings in marketing and advertising research. However, they seem incapable of attributing the precise sources to their findings, which prohibits an accurate and reliable literature review.

For a more detailed review of the prompts and LLM, see.

6. CASE STUDY 6: CODING OPEN ENDED

This case study offers an in-depth comparison and analysis of ChatGPT 4, Bard, and Claude AI in the context of coding open-ended survey questions. Using a recent survey question administered among 23 research companies that measure attention,[4] we first conducted thematic analysis and then created scatter plots to display the relationship between variables. The analysis provides valuable insights into the capabilities and limitations of these AI systems compared to human intelligence and offers guidance on effectively employing them in research and data analysis tasks.

Thematic Analysis

The prompt:

> Evaluate the responses and summarize into four or five segments.

Human analysis

- Attention as a function of **Visual Response**: Fixations can be regarded as the most granular elements of the attention stream. Companies may analyze gaze durations, reach, fixation intensity, and sequences.
- Attention as a **Biometric Response**: Such as heart rate variability or skin conductance. Reflects immersion, which is a measure of emotional connection and attention.
- Attention as a **Neurological Response:** Such as EEG. A comprehensive analysis of how consumers engage with information would be incomplete without considering emotion and memory in conjunction with attention.

[4] The question was: How does your company define attention?

- Attention as a **Digital Response:** Use observed data and combine it with tracking pixel data to train models that can predict human attention in-flight and at the point of transaction such as tracking pixels.
- **The Absence of Attention:** In layman's terms, inattention reflects not noticing something.

ChatGPT 4

- Attention as a measure of **emotional connection and engagement**: Several companies emphasize the emotional aspect of attention and view it as a measure of how engaged and involved individuals are with the content or stimulus.
- Attention as a **cognitive process**: Companies view attention as a cognitive process that involves the allocation of limited cognitive resources to process stimuli. They often highlight the different types and aspects of attention, such as sustained attention, orienting, top-down vs. bottom-up attention, etc.
- Attention as a **proxy for outcomes**: Some companies view attention as a proxy measurement for different outcomes, such as engagement, campaign effectiveness, and impression quality. They often incorporate multiple signals, such as visibility, situation, interaction, and physiological responses, to capture attention.
- Attention as a measure of **awareness and sensory processing**: Companies focus on attention as a measure of physical or mental orientation towards a stimulus, sensory processing, and presence in the viewing environment.

Bard

- **Segment 1: Focus and Prioritization**

  Audacy: Immersion

  MediaScience: Automatic or controlled focusing of limited cognitive resources

  Amplified Intelligence Technologies: When a human stops what they are doing, ignores other things around them, even if for a fleeting moment, to understand or learn something
- **Segment 2: Mental Engagement and Processing**

  Eye square: Meaningful attention

  NIQ-BASES: Mental prioritization of information

  Affectiva: Engagement

  Chilmark Digital: Power of content to stimulate individuals' ability and desire to process and respond
- **Segment 3: Sensory Processing and Awareness**

  Affectiva: Passive Attention

  Realeyes: Awareness to a stimuli while ignoring other stimuli
- **Segment 4: Measurement and Quantification**

  eye square: Fixations, gaze durations, reach, fixation intensity, sequences

  TVision: Person-level measure of attention at a second-by-second granularity

DoubleVerify: DV Authentic Attention¬Æ

Integral Ad Science: Visibility, Situation, and Interaction

Element Human: Visual orientation to the object of interest

MESH Experience: Experience Noticeability

Claude AI

- Attention as orientation/focusing of mental resources or senses towards a stimulus
- Attention as awareness and conscious processing of a stimulus
- Attention as a threshold for further information processing and memory encoding
- Attention tied to emotional engagement and arousal
- Attention as a predictor of downstream outcomes like ad recall and sales
- Distinctions made between passive attention/exposure vs active attention/engagement

Many companies mention the multidimensional nature of attention and that it involves both conscious/controlled and unconscious/automatic processes. Some also note the imperfection of attention as a metric and that it needs to be considered in conjunction with other metrics like emotion to get a full picture.

I don't have a strong assessment of which definition is "best" - they all provide value depending on the specific context and use case. But analyzing the different perspectives is helpful for me to better understand the nuances of how attention is viewed in the industry. Please feel free to provide any other sample data or clarify any of the definitions further. I'm happy to continue synthesizing and discussing how attention is characterized across different sources.

Following prompt:

> Can you extract 4 broad themes, name them and write a one sentence description of each.

Response:

1. **Orientation Attention** as the orientation of mental resources or sensory faculties towards a particular stimulus.
2. **Awareness Attention** as conscious awareness and active processing of a stimulus.
3. **Threshold Attention** as reaching a threshold of processing that enables encoding into memory and downstream impacts.
4. **Engagement Attention** that reflects emotional, motivational, or behavioral engagement with a stimulus.

While different terms were used to describe each feature of attention, most of the themes identified by all three LLMs could be mapped into a theme identified during the human analysis:

Human Analysis	ChatGPT 4	BARD	Claude
Visual response	Cognitive Process	Focus & Prioritization	Orientation
Biometric response	Awareness and Sensory Processing	Mental Engagement and Processing	Awareness
Neurological response	Emotional Connection and Engagement	Sensory Process and Awareness	Threshold
Digital response	Proxy for Outcomes	Measurement and Quantification	Engagement
Absence of attention			

Notably, human analysis predominantly focused on identifying differences in methodologies, in contrast to AI analysis which delved into broader conceptual aspects.

Upon reflection, AI analysis emerges as more insightful, capturing wider conceptual themes, while the human analysis seems more aligned with a systematic sorting of techniques, ironically akin to what one might expect from a computer's approach.

Scatter Plots:

A scatter plot is a type of data visualization used to display the relationship between two numerical variables. We wanted to see if each of the models were capable of doing at least a simple plot, that is, placing companies correctly on a two-by-two table. Upon visually examining the output of the three models, it appears that only ChatGPT 4 would categorize company missions in alignment with human analysis, specifically focusing on the words "attention" and "emotion". None of the models were capable of accurately performing data visualization, a task outside the typical strengths of Large Language Models (LLMs). However, with substantial training and integration of their APIs into a system, it's plausible that these models could contribute to data visualization systems in the future.

Human Analysis

In reviewing the mission statements of various companies, we categorized them based on their references to "attention" and/or "emotion" (see Figure 6). The majority claimed to study attention, while a smaller number indicated they measure emotion. This classification was strictly based on the presence of these specific terms in their mission statements.

Figure 6: Human Analysis Scatter Plot of "Attention" vs "Emotion"

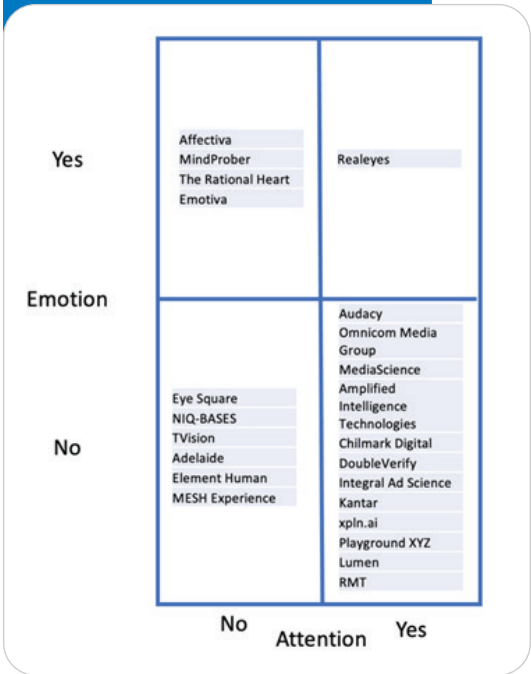

ChatGPT 4

- Classification: ChatGPT 4 successfully classified companies according to their mission statements, focusing on whether they mentioned "attention" or "emotion".
- Data Visualization: ChatGPT 4 struggled with data visualization, showing limitations in this aspect.
- Adaptability: It adapted to prompts for reclassification and improved in plotting companies in a 2x2 table, but still faced challenges in data visualization.

Figure 7: ChatGPT's First Scatter Plot of "Attention" vs "Emotion"

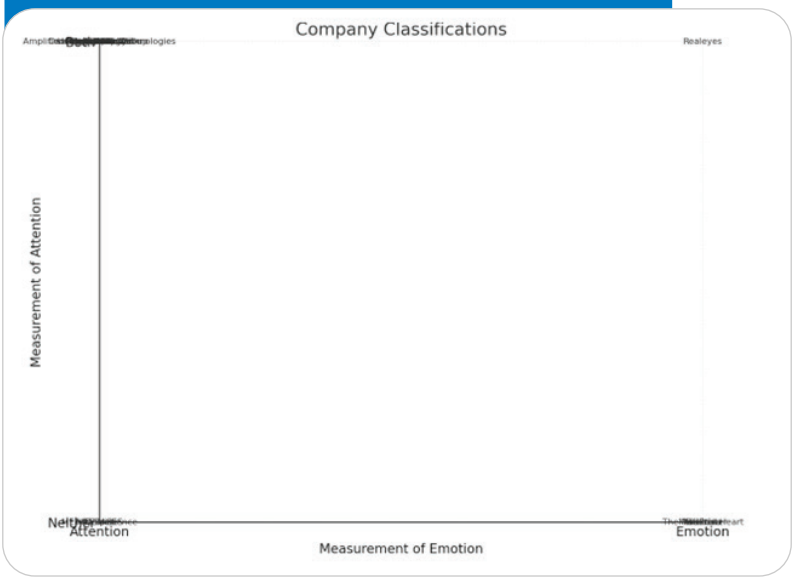

Figure 8: ChatGPT's Second Scatter Plot of "Attention" vs "Emotion"

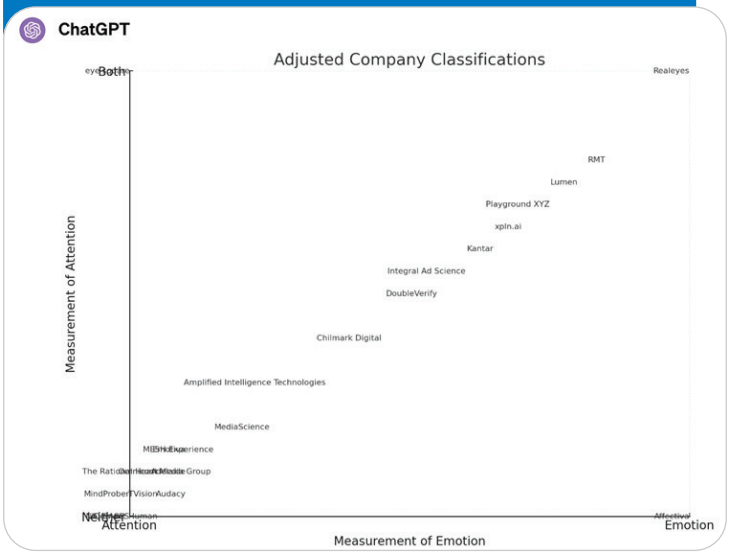

Bard

- Misclassification: Bard classified four out of 23 companies differently than ChatGPT 4 or human analysis. For example, Realeyes was classified under "Emotion" despite mentioning "attention" in its mission.
- Data Visualization: Bard was not effective in data visualization, with labels not properly placed in the correct columns. Bard can create a data visualization if receives a CVS file of data. Still, it's output is less satisfactory than ChatGPT 4's.
- Interpretation Variability: Bard's classification may reflect its interpretation of the mission statements, implying measurement of both "attention" and "emotion" even when not explicitly mentioned.

Figure 9: Bard's Scatter Plot of "Attention" vs "Emotion"

Measures Attention	Measures Emotion
Amplified Intelligence Technologies	Affectiva
Audacy	MindProber
Chilmark Digital	Realeyes
DoubleVerify	The Rational Heart
Integral Ad Science	Emotiva
Kantar (partially)	
Lumen	
MediaScience	
Omnicom Media Group	
Playground XYZ	

Claude AI

- Unusual Classification: Claude AI misclassified half of the companies when given the same prompt as the other models.
- Data Visualization: Although its classifications were incorrect, Claude AI successfully created a proper 2x2 table.
- Prompt Specificity: Claude AI did not improve in classifications even with prompts focused explicitly on the keywords “attention” and “emotion”.

Figure 10: Claude.AI’s Scatter Plot of “Attention” vs “Emotion”

	Attention	No Attention
Emotion	Kantar, MediaScience	Emotiva, The Rational Heart
No Emotion	Adelaide, Amplified Intelligence Technologies, DoubleVerify, Integral Ad Science, Lumen, Omnicom Media Group, Playground XYZ, Realeyes, TVision, xpln.ai	eye square, NIQ-BASES, Element Human, MESH Experience, Audacy, Chilmark Digital, Company Missions.docx, RMT, Affectiva, MindProber

Copy Retry

Key Takeaways:

All three LLMs responded with classifications that were more conceptual than the human classification which was more based on methods of measuring attention. It is worth clarifying that this study fundamentally was to evaluate the differences in methodology among the companies that measure attention. Hence, human classification might be biased by that objective. That said, if we ignore the purpose of the study, the LLMs may have done a better job of classifying the responses.

Only Claude AI did a reasonable and usable job of visually classifying the responses.

Task	ChatGPT 4	BARD	Claude.AI
Keyword Classification	Same as human	A few errors on a straightforward task	Very Poor
Data Visualization: 2 by 2 table	Some sense of the ask, could be trained	Didn’t understand the task	Best understanding of the task

For a more detailed review of the prompts and LLM responses, see,

7. CASE STUDY 7: REPORT DESIGN AND WRITING

In this case study, we examine AI's report writing capabilities. To do so, we use two different case studies: the ARF Privacy Report (September 2024) and the Erwin Ephron project[5] While the first is a deep investigation into how AI can be utilized for writing reports that are based on quantitative data analysis, the latter focuses on AI's capabilities to consolidate a book based on qualitative texts, including summarizing and reinterpreting newsletters, and organizing the bulk of texts into separate segmentations according to thematic logic. For both case studies, AI proves to be a useful tool for crafting the reports, despite obvious differences between the quantitative and qualitative data that informs each report. However, the process in which each report was crafted differ significantly, with each necessitating human guidance, involvement, and heavy editing and scrutiny, at different points throughout the process.

2023 Privacy Report:

For the Privacy Report, we conducted a comparative analysis of the process and the final outcomes between the report crafted by humans and the one generated by AI. The key finding is that the most effective strategy involved creating an outline and then providing the AI model with relevant tables in plain text (not in CSV format) to write sections of the outline. The AI-generated report was found to be impressively similar to the report created by a human. However, the AI models faced challenges when trying to write the report with the same inputs a human would use, such as the 2022 and 2023 Privacy Reports and Jupyter Notebook outputs of data analyses. The models struggled to parse input files and could only interpret data successfully when the analysis results were directly copied and pasted into the text box.

The models also faced difficulties in composing the entire report coherently. They would maintain relevance for a few paragraphs before beginning to "hallucinate," veering away from the original data they were provided with. Consequently, dividing the report into distinct sections and directing the AI to write each section individually is essential for maintaining focus and accuracy.

A major concern in employing AI for report writing emerged with the recurring issue of Claude AI and Bard fabricating data. These models frequently produced statistics that appeared plausible and were like those generated for the report, yet the figures were inaccurate. Since these are language models, it's crucial to meticulously verify all numbers and figures to ensure they are derived from actual data, rather than being artificially generated by the AI. In contrast, ChatGPT 4 proved to be an effective asset in report creation, provided that the initial organization of ideas is complete, and the data is already formatted for readability.

Additionally, both Bard and ChatGPT 4 have the capability to generate data visualizations given specific data inputs,[6] offering a time-efficient alternative to manually creating charts and graphs. Notably, ChatGPT 4 does this with a lot more ease and accuracy than Bard.

[5] For this case study only ChatGPT 4 was used.
[6] Bard, for instance, needs the data to be uploaded through in the format of a CSV file.

These findings highlight the need for converting data analyses into a format that AI can easily process and to thoroughly scrutinize anything an AI system does.

How the ARF Research Department Created the Report:

The process of composing the report began with a review of ARF Privacy Reports from previous years to understand their style and framework. Based on this we developed an outline, ensuring it included both recurring historical analyses and the outcomes of new questions introduced this year.

For every section, we examined the data to identify significant trends for emphasis and determined the most effective way to present these trends, creating data visualizations as needed to underscore key points. The writing phase involved weaving the relevant data analyses and visualizations into a cohesive narrative, providing necessary context or explanations for certain trends. Finally, we prepared a summary of the main trends for the report's introduction and crafted a conclusion discussing the broader implications of the privacy survey for the report's finale.

How AI Created the Report:

Ideally, an AI model would process the same inputs as a human, in this case the results from a data analysis in a Jupyter Notebook and reports from previous years, and generate a structured report inclusive of data visualizations. Nevertheless, our experience revealed that AI platforms necessitated substantial direction to yield practical outputs. We experimented with multiple approaches to steer the AI towards generating effective results, eventually identifying the most successful strategy.

	AI Input	AI Expected Output	Results
Attempt 1: Give AI the same information that a human would receive	• 2021 and 2022 Privacy Reports • Jupyter Notebook containing data analysis	Finalized version of the full 2023 Privacy Report	The first attempt involved giving AI the same amount of information that a human performing the task would receive. However, the models struggled to parse the Jupyter Notebook file format (despite all indicating that they could read these files), leading to time-out errors and falsified data. As a result, none of the models produced any usable output.

Attempt 2: Give AI CSV files rather than a Jupyter Notebook	• 2021 and 2022 Privacy Reports • 63 CSV files containing the results of the data analysis	Finalized version of the full 2023 Privacy Report	Because the models could not read the Jupyter Notebook file format, we then converted the output of each Jupyter Notebook cell into a CSV file and individually uploaded the files to each AI model. However, the large number of files caused the models to struggle with organizing the files into sections. The models could not recognize that certain tables might belong together (for example, Device Usage categorized by age and Device Usage categorized by gender). Bard returned hallucinated data only, while Claude.AI and ChatGPT were able to create some data visualizations based on the results, but could only create one or two at a time before requiring further human prompting. They are also unable to write a cohesive and fluid report, since they can only focus on a small section of the data at once. Thus, the volume of data in the analysis is an issue.
Attempt 3: Give AI pre-made data visualizations	• 2021 and 2022 Privacy Reports • All the human-generated data visualizations to be included in the 2023 Privacy Report	Finalized version of the full 2023 Privacy Report	The models could not extract meaning from the data visualizations. ChatGPT and Bard responded that they could not read images and asked the user to input written descriptions of each data visualization, at which point we stopped as this seemed counterintuitive. Claude.AI simply does not accept image files in any format. Thus, no output was produced from this method by any of the models.

Attempt 4: Give AI CSV files in plaintext format	• 2021 and 2022 Privacy Reports • Plain text versions of the CSV files for only the relevant analyses for the section in question	One section of the 2023 Privacy Report, such as "Devices and their Usage"	This attempt aimed to address the 2 primary issues that the models had faced thus far: 1. The models struggle to parse different file formats 2. The models struggle to read and organize the large amount of information in the analysis In response to (1), we inputted the contents of each analysis as raw text in CSV format directly into the chat box, rather than upload any files. In response to (2), we organized the analyses into pre-determined sections, and only gave the analyses to the AI models one section at a time, drastically reducing the amount of information AI received in one go. ChatGPT was able to produce a result that was nearly identical to something a human would produce, basically returning a final product. Claude.AI and Bard struggled more – they required additional prompting and the reports that they generated still contained large amounts of falsified information.

Optimizing AI for Data-Driven Report Writing: Strategies, Challenges, and Solutions:

- The most effective strategy for using AI to write a report from data analysis involved creating an outline and supplying the model with relevant tables in plain text format for writing individual sections.
- ChatGPT 4 generated reports remarkably similar to those created by humans when provided with these specific parameters.
- AI models were unable to write reports using the same inputs as humans (relying solely on the 2022 and 2023 Privacy Reports and the Jupyter Notebook output).
- The models struggled with parsing input files and could only interpret data successfully when it was directly copied and pasted into the text box, making data conversion time-consuming.

- AI models tended to stay accurate for a few paragraphs before deviating from the provided data, necessitating the division of the report into sections for AI to handle effectively.
- A major issue with AI report writing was the use of fabricated data by models like Claude AI and Bard, which often presented plausible but incorrect statistics.
- It's essential to double-check numbers and figures provided by AI models to ensure they are based on actual data.
- can be a valuable tool in report writing, particularly after ideas are organized and data is made AI-readable.
- Both Bard and ChatGPT 4 can create data visualizations from specific data inputs, offering a time-efficient alternative to manually crafting charts and graphs.

Wow: The Wit and Wisdom of Erwin Ephron

Between 1993 and 2010, Erwin Ephron, a highly influential figure in the field of advertising and media planning, published *The Ephron Letter*, over 150 newsletters that were for many in the industry a must-read for staying informed about the latest trends, theories, and discussions in the rapidly evolving world of media and advertising. In honor of his talents, The ARF published *Wow: The Wit and Wisdom of Erwin Ephron*, in January 2024 in tribute to Erwin's insightful analysis, commentary and critique. The full publication will be gifted to ARF WIDE (Workforce Initiative for Diversity and Excellence) contributors in early 2024.

Upon receiving the 150+ newsletters, we decided to utilize ChatGPT 4 to do the following:

1. Thematically organize the newsletters
2. Summarize each newsletter
3. Reinterpret each newsletter according to the current media landscape

The main challenges encountered in this project involved text uploading and accurate segmentation.

- It was not possible to upload the PDF document without retaining chat history, a concern for handling sensitive materials. However, for the Ephron project, this was not a critical issue since the newsletters were previously published.
- ChatGPT 4 had difficulty accurately segmenting the newsletters (150 newsletters in one PDF), failing to distinguish where one newsletter ended and another began, which was vital for summarizing and identifying themes.
- To address this, the PDF was manually converted into 150 separate Word documents, one for each newsletter.
- Another hurdle was the limit on uploading more than 10 documents at a time, making the process of uploading 150 documents laborious and time consuming.
- Once uploaded, ChatGPT 4 effectively summarized each text and reinterpreted each newsletter in the context of the current media landscape.
- For the book outline, ChatGPT 4 initially categorized the newsletters into sections based on suggested themes, which could be adjusted as needed.

- ChatGPT 4 was also asked to compare its outline with one created by the ARF, highlighting the strengths and weaknesses of each approach.
- Deciding on 10 sections, there were issues in distribution: ChatGPT 4 created duplicates and unevenly divided the newsletters among segments, leading to skewed distribution.
- The re-categorization attempts encountered several errors, resulting in a heavily skewed final distribution with many categories lacking newsletters.
- Ultimately, ChatGPT 4's suggested segmentation did not accurately reflect the content, requiring manual re-segmentation.

In sum, while ChatGPT 4 demonstrated effectiveness in summarizing individual newsletters and adapting their content to the current media landscape, it struggled with accurately segmenting a large collection of newsletters into coherent thematic sections. The AI encountered difficulties in processing a bulk PDF and required the content to be converted into individual Word documents. Even after this conversion, the AI's limitations in handling large-scale document organization became apparent, leading to issues like duplicate entries and uneven distribution across categories. This necessitated manual intervention for accurate categorization and segmentation of the newsletters for the project.

For a more detailed review of the prompts and LLM responses, see.

8. CASE STUDY 8: SYNTHETIC RESEARCH

Academic literature on using LLMs as substitutes for traditional advertising research is quite sparse. Currently, two primary methods to assess the viability of LLMs in advertising research are explored, offering insights into public opinion and consumer behavior:

1. Direct Questioning with LLMs: This approach involves posing specific questions to an LLM, relying on its training and search capabilities to generate responses that mirror published reviews, recommendations, surveys, press releases, and other sources, essentially replicating an actual research program. A caveat here is that online reviews might not be fully representative of the general population, similar to the limitations of online convenience surveys.
2. Creation of Synthetic Panels: This method involves creating a simulated panel of respondents that mirrors the attitudes of the target population, offering longitudinal consistency for tracking changes over time. When properly stratified, such synthetic panels can reflect the known demographic distributions of the relevant population. With an assumed 100% response rate, there would be no need for post-stratification adjustments.

Do LLMs Provide Estimates Similar to Those Collected Through ARF Primary Research?

In this case study, we explore whether LLMs can be a viable alternative for data collection, reflecting a broad spectrum of sources such as social media, reviews, and recommendations. Posing a series of questions similar to those used in the ARF's Sixth Annual Privacy Survey and DASH Survey, we evaluate how closely these AI systems could replicate the results of more quantitative research. The performance of ChatGPT 4 is quite impressive. Indeed, all three AIs provided conceptual and good analysis, showcasing their ability to understand and interpret advertising content.

Estimating Trust in Police and Media

Our prompt:

> What percent of American adults trust their local police?

Findings

- ChatGPT 4 yielded consistent findings with the ARF's Privacy Surveys over the last three years.

Our prompt:

> What percent of American adults trust the media in general?

Findings

- ChatGPT 4 yielded consistent findings with the ARF's Privacy Surveys over the last three years.
- Bard and Claude AI offered more comprehensive answers, reflecting the diverse estimates obtained from different surveys.

Media Behavior

Our prompt:

> What proportion of U.S. households get their television programming through broadband only?

Findings

- Bard and Claude AI provided answers that went beyond the initial query, offering a more thorough response.
- However, Claude AI cited older information, referencing a Nielsen report from 2021. Additionally, Claude AI's representation of the Nielsen data appeared to be inaccurate.

Our prompt:

> What percentage of Netflix subscribers share their accounts with friends or relatives outside of their household?

Findings

- The LLMs reported figures close to the ARF DASH report statistics, drawing from at least one source.

Creating Synthetic Survey Panels

Some scholars propose that using a synthetic panel of respondents can grant users more influence over the responses of LLMs. This can be achieved by tailoring responses to account for variations in demographics and regional attitudes. For instance, instructing an LLM to respond from a specific viewpoint, such as that of a researcher or marketer, can yield different answers. The concept of a synthetic panel enables users to simulate a range of controlled variables. To illustrate, we conducted a basic example of constructing a synthetic panel. Our findings show that all three AI models could straightforwardly generate synthetic survey panels, simulating respondents based on demographic data. You can task the LLM to assign synthetic respondent case ids and typical demographics for the regions, far more accurate control would be to provide estimates that captured all the interaction effects. So, you could task the LLMs to set the panels political preference to mimic the national distribution, answers about politics would be far more accurate if the panel was set by political preference by age by region of the country, …. . This approach demands meticulous attention to detail for each survey question to be effective – unless, of course, there's an emerging market dedicated to crafting high-quality synthetic panels.

For this test, each LLM was provided the marginal (not nested as recommended) demographic distributions for the US population. Then it was prompted:

> I would like you to create a synthetic survey panel of 300 U.S. adult respondents. They should each have an id ranging from 1 to 300. create the panel in such a way that they have a proper geographic distribution by 9 census regions, have a proper distribution by race/ethnicity, education, household size in which they live, whether they have cable television or broadband only or antenna.

All the LLMs could reasonably create such a panel with marginally accurate demographic distributions, although how the distributions were provided differed in that only ChatGPT 4 and Claude AI could ingest the data as an excel file.

Then each was asked:

> I would like to ask a survey question of each panelist: Are they Republican, Democrat or Independent. Answer the way you think they would, given their age, region and other demographic characteristics.

Results

ChatGPT 4

ChatGPT 4 initially reported everyone in a region as having the modal political preference of the region. For, example everyone from the South was assigned as a Republican. We then had to re-prompt:

> "Repeat the survey but assign a political party to each panelist Monte Carlo reflecting the proportion of Republicans, Democrats and Independents in that region. Therefore, not everyone from the South would be Republican."

While not ready for real analysis, this worked as each region had a varied distribution of political parties that reflected the distributions widely available on the internet.

At least for ChatGPT 4, using the API and with a significant amount of curation, a synthetic panel may be a way of improving LLMs as an alternative to conducting survey research.

Bard

Adeptly generated a synthetic panel and accurately represented the provided demographics. However, its estimations of political affiliations within the Monte Carlo model were excessively high. The political distributions heavily skewed towards Democrat affiliations, although the sum of the three affiliations (Democrat, Republican, and Independent) did correctly total 100% for each region.

When asked to display a bar chart showing the demographic distributions of the panel Bard required more complex data ingestion processes and lacked graphic visualization capabilities.

Claude AI

Claude AI provided the most accurate and direct simulated political affiliations without needing re-prompting. Overall, the distribution of the political parties is consistent with published estimates.

Claude AI doesn't do any form of visualization (bar chart) and its CSV format is textual and difficult to export without coding.

Key Takeaways:

TASK	ChatGPT 4	BARD	Claude.AI
Estimating Trust and Media	Reasonable estimates	Reasonable estimates	Reasonable estimates
Visualization of Data	Can present bar charts	No graphic visualization	No graphic visualization
Generating Synthetic Panel	Straightforward	Straightforward	Straightforward
Output of Synthetic Cases	Provided link to Excel file	Poor, CSV as text within the response	Poor, CSV as text within response
Panel Political Response % By Party	Had to re-prompt to train Monte Carlo, initially each region assigned 100% to modal party. Reasonable Republican distribution by region.	Had to re-prompt to train Monte Carlo, initially each region assigned 100% to modal party. Inaccurate Republican distribution by region.	Accurate political representation without re-prompting.

ChatGPT 4 was the easiest to use due to its superior input and output abilities.

Claude AI produced the most reasonable estimates for political affiliation, aligning with published data.

While the AI models can simulate research processes, they will require careful curation and control of variables to provide an alternative to survey research.

This section was only intended to illustrate how to go about creating a synthetic panel and not to judge how accurate they would be. We spent at best a day experimenting with different prompts at a marginal level of demographic control – clearly insufficient to mimic the US adult population. Monte Carlo was one solution to an obvious failure. A real development program likely will have to catalogue all the failures and find solutions for the many kinds of failures that are likely to exist.

For a more detailed review of the prompts and LLM responses, see.

CHAPTER 5: ETHICS AND TRANSPARENCY IN AI-DRIVEN ADVERTISING RESEARCH

Having illustrated the kinds of actions people might take with LLM, we now address the ethical issues related to those kinds of advertising research functions.

Advancements in Artificial Intelligence (AI) have sparked significant interest within both the media and the public. As AI systems, including robots, chatbots, avatars, and other intelligent agents, transition from mere tools to autonomous entities and collaborators, there is a critical emphasis on researching and comprehending the ethical implications associated with these systems. This shift prompts a series of important inquiries: What constitutes decision-making in the context of AI systems? What ethical, societal, and legal repercussions arise from their actions and decisions? Can AI systems be held accountable for their conduct? How can we exercise control over these systems as their learning capabilities lead them into states that may only tangentially relate to their initial design and setup? Should we permit such autonomous innovation in commercial systems, and if so, how should their use and development be regulated? These questions, along with numerous others, currently occupy the forefront of attention. The way society and our institutions address these inquiries will significantly impact our trust levels, the broader influence of AI in society, and ultimately, the existence of AI itself (Boddington, 2017; Bostrom & Yudkowsky, 2018; Dignum, 2018).

The evolution of business models, sales processes, customer service options, and marketing information systems (Donthu & Gustafsson, 2020), necessitates a careful consideration of ethical issues and data protection concerns (Ameen et al., 2020a; Etzioni & Etzioni, 2017). For instance, the collection of data through speech recognition, including clients' tone of voice when interacting with voice bots, to enhance marketing strategies, must adhere to the General Data Protection Regulation and obtain client consent (Butterworth, 2018) (See the future of Voice and Conversational AI in the next chapter. To mitigate consumer skepticism and prevent speciesism towards AI, practitioners should uphold ethical standards (Stone et al., 2020) and prioritize data protection (Kolbjørnsrud et al., 2017). Understanding privacy in this context involves recognizing its varied cultural and regional interpretations. An ethical quandary arises when considering the collection of user data without explicit consent, prompting a need for methods that aggregate and anonymize data to safeguard individual identities. The importance of secure data storage solutions cannot be overstated, with regular audits and updates being crucial to prevent data breaches. In terms of regulations and compliance, businesses must navigate the complexities of global data protection laws like the GDPR and CCPA, understanding that non-compliance can have serious repercussions. Building trust with users is paramount, achieved by informing them about the use of their data and offering options to opt-out or limit data collection (see e.g., Ipsos, 2023).

Although recent advancements in information technology and AI are facilitating better coordination and integration between humans and technology, developing what has been termed Human-Aware AI, it is still far from being able to function as a "team member," adapting to the cognitive strengths and weaknesses of human collaborators (Korteling et al., 2021). To this extent, while AI models can make highly accurate predictions, they do so in ways that are difficult for humans to comprehend. This lack of transparency can be problematic in various domains, including healthcare, finance, and autonomous vehicles, where it's crucial to understand why a particular decision was made. The growing complexity of AI and machine learning models, such as deep neural networks, which are often considered "black box" systems strengthen the need for explainable AI (XAI).

XAI refers to the concept of designing and developing artificial intelligence systems and machine learning models in a way that makes their decisions, predictions, and reasoning

processes understandable and interpretable by humans. In other words, XAI seeks to provide insights into how AI systems arrive at their conclusions, allowing users to grasp the rationale behind those decisions.

Explainable AI techniques aim to address this issue by providing *Transparency* and visibility into the inner workings of AI models;[7] *Interpretability*, namely human-understandable explanations for AI predictions (in the form of visualizations, textual descriptions, or other formats that make it easier for users to grasp the reasoning behind AI decisions); *Accountability*, namely allowing for the identification of biases, errors, or ethical concerns in AI models.

Explainable AI is particularly important in applications where trust, fairness, and safety are critical. It can help AI practitioners, regulators, and end-users have confidence in AI systems and ensure that they align with ethical and legal standards. XAI techniques continue to evolve and play a crucial role in the responsible deployment of AI technologies across various industries (Goebel et al., 2018; Gunning et al., 2019; Xu et al., 2019).

Another way to mitigate or counter AI is through systems that can differentiate between content created by AI and that made by humans (GPTZero, Fictitious.ai, and Writer.com as examples). This emerging requirement addresses concerns about authenticity and origin in various digital media, ensuring clarity and credibility in a landscape where AI-generated content is becoming more common and sophisticated (Fried, 2023).

The surge of interest in AI capabilities has also given rise to numerous inquiries concerning the societal repercussions, potential misuse, risks, and governance of these innovations, all of which hold paramount significance. One of the major concerns is about the potential misuse of AI platforms in spreading disinformation (Metz, 2023) and the challenges in identifying AI-generated content (Spitale et al., 2023). Up until recently ChatGPT 4 lacked internet access and possessed only limited knowledge of global developments post-2021 (Stokel-Walker & van Noorden, 2023). As the sphere of knowledge is continuously evolving, this constraint occasionally led to the delivery of outdated or erroneous responses. At the time of writing this report, ChatGPT 4 uses Microsoft's Bing search engine to locate relevant information when producing responses to prompts. As of the writing of this report, ChatGPT 4 4 also claims to have been updated to March 2023. To be sure, while it often provides relevant and useful results, its accuracy can vary. A related problem is the tendency to generate seemingly credible but ultimately fictitious citations that lack real-world sources when prompted to incorporate current references (Choi et al., 2023).

An over-reliance on AI may have negative consequences in terms of a decline in higher-order cognitive skills such as creativity, critical thinking, reasoning, and problem-solving (Farrokhnia et al., 2023).

Another concern is about AI's ability to discern between authentic and fabricated content, including distinguishing its own outputs from those made by humans. This issue underlines

[7] Transparency is not only about the AI models but also about data usage, necessitating clear communication about the type of data collected and its purpose. Stakeholder engagement is crucial, involving consumers, regulators, and others in discussions about AI transparency and collaborative efforts to set industry standards. Moreover, educating the public through workshops, webinars, and educational content is essential for demystifying AI in advertising research, fostering a broader understanding and acceptance.

the potential risks in trusting AI-generated information and highlights the ease with which malicious actors could exploit such tools to generate large volumes of deceptive content (Coldewey & Lardinois, 2023). In the context of advertising and, occasionally, advertising research, the ethical implications of these AI-related issues are significant. The potential decline in cognitive skills like creativity and critical thinking due to AI reliance raises questions about the quality and originality of advertising content. Furthermore, AI's challenge in distinguishing between authentic and fabricated content can lead to ethical dilemmas in advertising practices, where discerning the truthfulness of AI-generated information becomes crucial.

Bias and fairness in AI models present another significant challenge. Bias can creep into AI models through various channels, and distinguishing between model bias and data bias is essential. In advertising research, the consequences of bias can lead to the misrepresentation or exclusion of certain groups, fostering negative brand perceptions. Techniques like adversarial testing are employed to detect bias in AI models, with data augmentation and diversification being vital for creating balanced datasets. Continuous monitoring is key to ensuring AI models remain fair, adapting as societal norms and values evolve.

Relatedly, there has been a growing focus on the targeted applications of AI in the realm of social well-being. This domain has attracted a multitude of stakeholders, including charitable organizations such as DataKind (established in 2012), academic initiatives like the Data Science for Social Good (DSSG) program at the University of Chicago (established in 2013), and international entities like the UN Global Pulse Labs. Additionally, there has been a proliferation of AI for Social Good workshops at prominent conferences, including the 2018 and 2019 NeurIPS conferences, the 2019 ICML conference, and the 2019 ICLR conference. Corporate support in this endeavor has been evident through initiatives such as Google AI for Good Grants, Microsoft AI for Humanity, Mastercard Center for Inclusive Growth, and the Rockefeller Foundation's Data Science for Social Impact, among others (Tomašev et al., 2020). Recent studies have demonstrated the potential benefits of harnessing AI for societal betterment. For instance, Amnesty International and ElementAI demonstrated how AI can assist human moderators in identifying and quantifying online abuse against women on Twitter. Anticipated enhancements in both data infrastructure and AI technology hold the promise of enabling an even wider array of potential applications for AI in the service of societal good.

Conclusion:

While AI-driven advertising research offers transformative potential, it also presents ethical challenges. As the industry advances, balancing innovation with ethical considerations becomes paramount. Companies that prioritize privacy, fairness, and transparency not only adhere to regulatory standards but also foster trust with their audiences, paving the way for a more responsible and inclusive AI-driven advertising landscape.

CHAPTER 6: THE FUTURE OF AI IN ADVERTISING RESEARCH

1. TRENDS TO WATCH

Hyper-Personalization:

The move from broad audience targeting to individualized messaging. AI will allow advertisers to create bespoke content tailored to individual preferences, behaviors, and real-time circumstances. This technology can allow advertisers to create highly personalized content, tailored to the unique preferences, behaviors, and even real-time circumstances of each individual. By leveraging AI, advertisers can analyze vast amounts of data to understand and predict consumer needs and preferences at an individual level. This enables the crafting of bespoke messages and content that resonate more effectively with each consumer, potentially increasing engagement and effectiveness of advertising campaigns (Deveau et al., 2023). At the same time, browser-imposed privacy restrictions, like those from Safari and Chrome, affect personalization algorithms creating the need to find solutions to mitigate the adverse effects of privacy restrictions and necessitating finding ways to balance user privacy with the efficacy of personalized recommendations (Korganbekova & Zuber, 2023).

Augmented Reality (AR) and Virtual Reality (VR) Advertising:

Leveraging AI in advertising within augmented reality (AR) and virtual reality (VR) spaces creates immersive and interactive experiences. These AI-driven ads can dynamically adapt and respond to user interactions, providing a highly personalized and engaging advertising experience. In these environments, advertisements are no longer static but become interactive elements that users can engage with in real-time, enhancing the impact and memorability of the marketing message. This approach represents a significant advancement in how brands can connect with their audience, offering novel and engaging ways to experience advertisements (ARF Immersive Advertising Study, 2023). Bayrak et al. (2020) focused on consumer attitudes towards augmented reality advertising, revealing differences between Turkish and German consumers towards AR ads. Kitsopoulou & Lappas (2023) review 22 studies on the effectiveness of AR/VR advertising for product and services promotion to show how integrating AR/VR in advertising can foster more positive attitudes towards both the advertising itself and the brands being advertised, compared to traditional advertising methods. This enhanced advertising approach positively impacts consumer behaviors, including increased purchase intentions, better brand recall, and improved perceptions of ad credibility. AR/VR ads are particularly effective due to their unique features like 3D effects, innovation, vividness, interactivity, personalization, telepresence, realism, and the capability for real-time interaction.

Voice and Conversational AI:

As voice search becomes more prevalent and virtual assistants like Alexa and Siri are increasingly integrated into daily life, there is a rising trend in voice-activated advertising

(Widjaya, 2023). This shift means that advertisers are likely to develop strategies that leverage voice commands and responses. Such advertising could involve interactive voice ads that engage users in a conversation, provide personalized responses, or guide them through various options and choices based on their verbal input (for e.g., Pizza Hut and Estée Lauder's implementation of voice-activated ads in their marketing campaigns – Parachuk, 2021). The growth of voice-activated technology paves the way for innovative advertising approaches that cater to this hands-free, voice-driven user experience.

With voice searches being more conversational in nature, there's a growing importance for long-tail keywords, question-based queries, and local Search Engine Optimization (SEO). Additionally, securing featured snippets on search engines becomes vital for visibility in voice searches. Voice search presents opportunities for brands to build trust by providing accurate and quick responses to queries, personalize user experiences, and integrate with e-commerce for seamless shopping experiences. Brands can also customize voice assistants to align with their identity, offering unique branding opportunities.

However, challenges accompany these opportunities. The rise of voice search might reduce screen time and website visits, as users receive direct answers from voice searches. The accuracy of voice recognition across different accents and languages is another challenge. Addressing privacy concerns and innovating new monetization strategies that fit a voice-first world are also critical (Widjaya, 2023).

Predictive Analytics:

Advanced AI models in marketing and consumer behavior analytics have progressed beyond just analyzing past consumer behaviors. They now possess the capability to predict future actions, which is a game-changer for brands in terms of strategy and consumer engagement. These models use historical data, including past interactions, purchases, and engagement patterns, to forecast future consumer actions. This predictive capability enables a shift from reactive to proactive strategies in marketing, allowing brands to anticipate consumer needs and preferences before they are explicitly expressed.

The use of predictive analytics facilitates personalization at scale, allowing brands to tailor their marketing efforts much more precisely to individual consumer profiles (Arora et al., 2008). It can enhance consumer experience by tailoring experiences, offers, and communications effectively (Aggarwal & McGill, 2007). Another application of predictive analytics is for dynamic marketing campaigns adjusted in real-time based on AI's ongoing learning about consumer behaviors to ensure that marketing efforts are continuously optimized for maximum effectiveness (Li & Kannan, 2014). Demand forecasting and identifying emerging trends can also be enhanced through predictive analytics. For the first, by predicting future buying patterns, brands can better manage inventory, supply chain logistics, and production planning. For the latter, by analyzing broad consumer behavior patterns, AI can help brands stay ahead in product development, service offerings, and market positioning.

Emotion Detection and Response:

Analyzing facial expressions, voice intonations, and other behavioral cues through AI is transforming the landscape of targeted advertising (see ARF Attention Measurement

Validation Initiative: Phase 1; ARF Attention Measurement Validation Initiative: Literature Review). This advanced approach leverages the capabilities of AI technologies to interpret human emotions, paving the way for more personalized advertising content. Facial Expression Analysis is one area where AI can be used: AI algorithms, trained in facial recognition technologies, can discern various emotions from facial expressions (Lewinski et al., 2014). By interpreting emotions like happiness or disappointment, AI provides valuable insights into consumer reactions to specific advertising content. Similarly, AI's ability to analyze voice intonations is crucial for emotional gauging. Variations in speech elements can be indicative of different emotional states and voice analysis can now accurately assess emotional states, such as excitement or frustration, which are key in tailoring advertising strategies. Beyond facial and vocal cues, AI systems are also adept at interpreting other behavioral indicators, including body language and eye movements. These cues offer a more comprehensive understanding of a user's emotional state and engagement level (Baltrušaitis et al., 2016).

With these insights, AI enables advertisers to customize their content more effectively. Positive emotional responses can guide the development of similar future content, while negative responses can prompt adjustments. This dynamic approach ensures that advertising is not only effective but also resonates with the audience's emotional preferences. This technology goes beyond benefiting advertisers; it significantly enhances user experience. As ads become more aligned with users' emotional states and preferences, they become more engaging and relevant. Furthermore, as users' emotional responses are detected, the AI can instantly modify the content, making the ads more responsive and interactive. For instance, facial recognition advertising involves utilizing sensors that can identify a customer's face, subsequently altering the advertisement's display in real-time. This approach aims to develop adaptive and dynamic advertisements that modify their content to align with a person's interests as soon as they engage with the ad (Kuligowski, 2023).

However, this technological advancement is not without its challenges. The use of AI for emotion analysis in advertising raises significant ethical and privacy concerns. Users might find the idea of emotional monitoring invasive. Responsible use of this technology, with adherence to privacy laws and clear communication about data usage, is crucial.

Quantum Computing:

The potential integration of quantum computing into AI could revolutionize the field, offering unprecedented advancements in data processing speeds and the complexity of AI models. Quantum computing's use of qubits[8] in principle enables it to perform numerous calculations simultaneously, potentially offering exponential increases in data processing speeds compared to classical computing. A simple analogy follows: we all know that the digital revolution is based on a bit being turned on or off – the state is 1 or 0. In the quantum world, this state can be any superposition – any linear combination of "on" or "off" (the Schrodinger cat state). Furthermore, a quantum device allows simultaneous operations on many qubits, as they're "entangled." This can increase the representative power by, well, a whole lot. This rapid

[8] In classical computing, data is encoded in binary bits, which can be either 0 or 1. In contrast, qubits, the basic units of information in quantum computing, operate according to the principles of quantum mechanics, offering a more complex and powerful way of processing information. Rather than the infinite positions of electrons, the quantum world is categorized by infinite possible linear combinations of on/off states. The increase in computational power is due both to this, and the fact that these qubits all talk to each other and can operate at the same time.

processing capability could be highly useful in marketing, where analyzing vast amounts of consumer data for insights is essential. Similarly, the ability of quantum computing to handle complex, large datasets could revolutionize consumer behavior analysis. This means more sophisticated AI models in marketing that can analyze consumer patterns and preferences with unprecedented depth and accuracy. Finally, quantum computing could significantly shorten the time required for training complex predictive models in marketing. This rapid training capability means quicker adaptation to market trends and consumer behavior shifts, leading to more responsive and effective marketing strategies.

In terms of digital ad targeting, unlike traditional print advertising, digital ads leveraging quantum technology can deliver more personalized and successful campaigns. This is particularly relevant as quantum computers can analyze data more efficiently, enabling marketers to better understand target audiences and create more effective advertisements. Quantum computing could offer novel ways to track ad campaign success, especially as traditional methods like cookies face challenges from browser restrictions. Its computational power could fill the gaps left by conventional tracking methods, aiding in more precise ad targeting. In a landscape where consumers expect brands to create content, quantum computing can aid in producing unique content more efficiently. Quantum machine learning (QML) algorithms can generate creative content swiftly, augmenting the capabilities of AI in content creation (Hughes-Castleberry, 2023).

2. CHALLENGES AHEAD

Data Privacy and Regulation:

As AI becomes more pervasive, so does the scrutiny on data collection practices. Complying with evolving data privacy laws will be paramount.

On April 21, 2021, the European Commission proposed the AI Regulation, the first comprehensive legal framework for AI to encourage investment and innovation while ensuring safety and fundamental rights. This draft regulation proposes harmonized rules, global reach, turnover-based fines, and emphasizes transparency, risk management, and accountability. It covers the entire AI lifecycle and applies to providers, users, distributors, importers, and resellers, including those outside the EU if their AI systems operate within it. The regulation categorizes AI practices into three tiers: unacceptable, high-risk, and low-risk. Prohibited practices include AI for social scoring, large-scale surveillance, and adverse behavioral influencing. High-risk AI systems, like those in critical infrastructure or justice, are allowed under strict controls. Low-risk AI systems are subject to basic transparency requirements. The regulation's impact on market research is minimal, but it reinforces existing self-regulatory regimes and emphasizes careful management of biometric data and other sensitive areas. On December 9, 2023, a tentative agreement was reached by Parliament and Council representatives on the Artificial Intelligence Act.

In the US, AI regulation took center stage leading to President's Biden's October 2023 executive order which called for increased transparency and new standards in AI. This order laid the foundation for a US-centric AI policy, emphasizing industry-friendly best practices and allowing various agencies to develop their own sector-specific regulations. The 2023

legislative session in the U.S. has seen an unprecedented number of state AI laws being proposed, eclipsing previous years. Ten states have incorporated AI regulations within broader consumer privacy laws set to take effect or already in place this year. Additionally, numerous states have introduced bills with similar objectives. Various states are establishing task forces to scrutinize AI's role in sectors like healthcare, insurance, and employment. A notable example is a law in New York City, part of the broader consumer privacy legislation, which specifically addresses AI in hiring practices and has garnered national interest. Apart from the California AI-ware Act, which governs the use of generative AI in government applications, several other bills are aimed at addressing the potential harms of generative AI. These legislative efforts largely concentrate on mitigating the issues arising from AI-generated images and videos (Zhu, 2023).

According to the MIT Technology Review (Ryan-Mosley et al., 2024) we can anticipate a risk-based regulatory approach to AI, similar to the EU's AI Act, where AI types and applications are assessed based on their risk levels. The National Institute of Standards and Technology has proposed this framework, which is now set to be implemented across various sectors and agencies.

Ethical Concerns:

Beyond data privacy, the ethical implications of AI decisions, biased algorithms, and the potential misuse of hyper-personalized content will be areas of concern.

Market research has long recognized the right of individuals to control their personal data, guided by transparency in data collection, protection of personal data, and ethical behavior. With AI and secondary data use challenging traditional methods, there's a need for a more outward-focused ethical framework, extending beyond protecting participant and client interests to actively doing good.

This shift, aligning with civil society and legislative expectations, calls for the market research industry to lead in setting behavioral standards. For AI in market research, ESOMAR (Cooke & Passingham, 2022) propose an ethical framework based on Floridi & Cowls (2019) Framework that emphasizes beneficence, non-maleficence, autonomy, justice, and explicability. This framework, influential in AI4People and the European Commission's Ethics Guidelines for Trustworthy AI, can offer a starting point to develop AI guidelines, considering diverse perspectives. As AI evolves in market research, adapting to this ethical approach is important for guiding industry standards and responding to the global regulatory environment.

Technical Complexities:

As AI models grow more sophisticated, so will the need for advanced expertise and infrastructure.

The dynamic nature of AI technology means that the learning curve in the advertising industry will become steeper and continuous. Professionals will need to stay abreast of the latest AI trends, tools, and methodologies to remain relevant and effective in their roles. To help with this, educational institutions and industry bodies will need to offer targeted training and up-skilling programs. These initiatives should focus not only on technical skills related to AI but also on enhancing creative, analytical, and strategic thinking abilities. However, other than the

more technical coding examples needed to evaluate survey data, the case studies presented in Chapter 4, could be experimented with, played with, explored with anyone with no coding experience. There are few barriers to entry introducing LLMs to an advertising research project.

Companies in the advertising sector will need to manage the transition effectively. This includes providing support for employees undergoing re-skilling, fostering a culture of lifelong learning, and redesigning job roles to accommodate the coexistence of AI and human expertise. Since the future of advertising is not about AI replacing humans, but rather about humans working in tandem with AI, developing skills to collaborate effectively with AI, understanding its capabilities and limitations, and using AI as a tool to enhance human-driven strategies are all of ultimate importance.

Public Perception:

Balancing innovation with concerns about the "creepiness" factor. Brands will need to be cautious about not overstepping boundaries and making users feel overly monitored. This balance is crucial for brands that want to leverage advanced technologies to personalize and enhance customer experiences without alienating their audience. The sense of discomfort or unease that consumers may feel when they perceive that a brand has too much insight into their personal lives or behaviors can occur when personalized marketing efforts become too intrusive or intimate, giving the impression that the consumer's privacy has been violated. That said, there are generations that don't change a channel by a remote, but instead talk to their television – and don't find this creepy, even though their television is an advertising platform harvesting data about their preferences.

Boerman & Smit (2023) conduct a systematic review of 84 articles to reveal three main contexts in which privacy is a key theme in advertising: as part of the ethical and regulatory considerations of advertising, in relation to personal characteristics that vary among consumers, and as a factor influencing how consumers respond to and are affected by advertising (for e.g., Lina & Setiyanto, 2021). Looking forward, Boerman and Smit (2023) address the growing use of personalized advertising in public spaces, privacy fatigue (also termed privacy cynicism), and ways in which to deal with constraints to personalization.

In balancing personalization and privacy, brands should focus on engaging consumers effectively without invading their privacy. They must ethically manage consumer data, complying with regulations like GDPR and CCPA. Building consumer trust is crucial, requiring transparency in data usage and respect for privacy. Brands should avoid overreliance on technology for data analysis and consider ethical implications in marketing strategies. Regular consumer feedback helps in adjusting strategies. The use of AI and analytics must be balanced with privacy concerns. The goal is to enhance consumer experiences innovatively without overstepping privacy boundaries.

Economic Implications:

Potential job displacements in the advertising industry, requiring re-skilling and up-skilling initiatives.

The advent of AI in the advertising industry brings with it a significant shift in the nature of work and the skill sets required. This technological evolution could potentially lead to job

displacements, as AI systems and algorithms become capable of performing tasks that were traditionally done by humans. These changes necessitate a focus on re-skilling and up-skilling initiatives to prepare the workforce for the new landscape. Since AI is exceptionally adept at automating routine, repetitive tasks, in advertising, this could mean automation in areas like data analysis, customer segmentation, and even some aspects of creative design. As AI takes over certain tasks, new roles and skills will potentially emerge. For instance, there will be a growing need for AI trainers who can teach AI systems how to mimic human-like decisions in advertising contexts. Similarly, roles centered around AI ethics, compliance, and interpretation of AI-driven insights will become crucial. At the same time, while AI can handle data-driven tasks efficiently, it is clear that human creativity and strategic thinking are not easily replicated by machines. Professionals in the advertising industry may need to pivot more towards roles that leverage these uniquely human skills, such as creative direction, strategy development, and emotional engagement in advertising campaigns.

While numerous startups have begun to offer AI marketing services and businesses and universities started offering AI marketing certifications, major brands have not yet restructured their organizations or created AI-specific leadership roles. Indeed, the absence of roles like VP of AI marketing is noticeably apparent, making us question the gap between AI's perceived impact and its actual implementation in marketing leadership. The Wall Street journal (Coffee, 2023) noted that as of November 2023, job listings mentioning AI in marketing were 8% lower than the previous year, despite the rise of AI startups like OpenAI. This is in contrast to *Indeed*, which saved $10 million using generative AI for content development in 2023. However, AI's influence on marketing seems less significant compared to other fields – for instance, sales job listings were almost three times more likely to mention AI compared to those in marketing.

Some B2B companies and advertising firms like WPP have introduced roles such as chief AI officer and head of AI, focusing mainly on promoting products and services to business clients. Apart from Coca-Cola, who promoted two executives to the newly created roles of global head of generative AI and global head of marketing AI, such titles remain relatively rare in the broader marketing industry.

Additional economic implications can be opening new opportunities for revenue generation marketing. For instance, AI-driven insights can lead to the development of new products or services, targeted advertising, and dynamic pricing strategies. Additionally, AI's ability to analyze and predict consumer behavior with high accuracy can lead to more effective allocation of marketing budgets, ensuring higher returns on investment. AI tools can potentially level the playing field, allowing smaller businesses to compete more effectively with larger players by providing insights and automation that were previously only available to those with extensive resources (Korganbekova & Zuber, 2023) demonstrate that their probabilistic recognition algorithm can improve visibility and revenue for smaller sellers. This approach effectively counters the disproportionate effects of privacy restrictions on vulnerable consumer groups and smaller sellers) . At the same time, with the increased use of AI in advertising, companies will need to invest more in data management and security. Ensuring data privacy and complying with regulations like GDPR can entail significant costs.

3. OPPORTUNITIES FOR RESEARCHERS AND BRANDS

New Data Sources:

The integration of wearables, IoT (Internet of Things) devices, and smart home technology into the advertising ecosystem is set to provide a wealth of new data points that can significantly enhance the way advertisers analyze consumer behavior and tailor their strategies. Each of these technologies contributes uniquely to the pool of consumer data. Wearables, including smartwatches, fitness trackers, and even smart clothing, are continuously collecting data about users' health, activities, and preferences. This data includes heart rates, exercise patterns, sleep quality, and sometimes even location data. For advertisers, this information can be invaluable in understanding consumer lifestyles and habits, allowing for more personalized and targeted advertising. IoT Devices – the Internet of Things encompasses a wide range of devices connected to the internet, from smart refrigerators and thermostats to connected cars and security systems. These devices provide real-time data on consumer preferences and behaviors within the home and beyond. For instance, a smart refrigerator could provide insights into a family's eating habits, which could be used to tailor grocery or meal delivery service ads. Smart Homes – the broader ecosystem of a smart home, which might include IoT devices like smart lighting, voice assistants, and automated cleaning systems, provides a holistic view of a consumer's lifestyle choices and preferences. This comprehensive data can help advertisers create a more complete profile of a consumer's daily life, routines, and potential needs.

The data from these sources offer several advantages for advertisers, including *Hyper-Personalization*, *predictive modeling* that enables advertisers to anticipate consumer needs and offer products or services at the most opportune times, *enhanced engagement* as ads are more relevant to the consumer's current lifestyle and needs, and *new advertising channels*, from push notifications on a smartwatch to voice ads through a smart speaker.

Collaborative AI:

The integration of AI tools in advertising research emphasizes a collaborative model where AI enhances rather than replaces human creativity and insight. This approach is particularly relevant in the advertising industry, where creativity and a deep understanding of human behavior are crucial. AI tools that work alongside human researchers, enhancing creativity and insight rather than replacing them (Girotra et al., 2023). AI can assist advertising researchers by processing large-scale consumer data, identifying emerging trends, and performing predictive analytics. For instance, AI's ability to analyze extensive consumer data can reveal patterns and preferences that might be missed by traditional research methods. Human researchers can interpret these insights, integrating them with contextual knowledge to develop advertising campaigns that resonate more deeply with audiences. Additionally, AI can support the creative process in advertising by offering data-driven insights, generating initial creative concepts, and proposing design variations (Ipsos, 2023). These AI-generated ideas serve as a foundation for human researchers, who can then refine and enhance them with their creative expertise and understanding of the target audience. As such, this partnership allows researchers to focus on higher-level creative and strategic aspects, leveraging AI's insights to inform more effective advertising strategies.

In a collaborative setting, researchers can ensure that AI's data-driven insights are balanced with an understanding of the emotional and cultural factors that drive consumer behavior. This collaborative model fosters a dynamic learning environment where AI systems improve through human feedback, and researchers adapt to new data-driven methodologies. This synergy ensures that advertising research remains innovative and responsive to evolving consumer landscapes.

Sustainability and Ethical Branding:

Using AI in advertising research to analyze and respond to consumer demands for sustainability and ethical practices is increasingly relevant as consumers become more environmentally and socially conscious. AI can analyze vast datasets from various sources – social media, consumer surveys, online forums – to gauge public sentiment about sustainability and ethics. This analysis helps in understanding the evolving expectations of consumers regarding environmental responsibility and ethical business practices. AI can also be instrumental in measuring the impact of a brand's sustainability initiatives, providing tangible data on consumer engagement, campaign effectiveness, and changes in consumer perception. This measurement is key to understanding the ROI of sustainability-focused advertising and guiding future strategies. Finally, when using AI for advertising research, it's essential to consider ethical implications, particularly in terms of data privacy and the potential biases in AI algorithms. Transparency in how consumer data is used and ensuring that AI models are free from biases are critical aspects of using AI in advertising research. Using AI in advertising research offers a path to more responsive, responsible, and effective research. It reflects a growing recognition in the advertising industry of the need to align with broader societal values and consumer expectations, using technology as a tool for greater engagement and impact.

In sum, the future of AI in advertising research is brimming with potential. As we venture into this new era, the blend of opportunities and challenges will reshape the landscape. By keeping abreast of emerging trends, anticipating challenges, and seizing new avenues for exploration, researchers and brands can pioneer a future where AI not only amplifies advertising research efforts but also deepens the connection between brands and their audiences.

CONCLUSION

1. EMBRACING THE AI REVOLUTION IN ADVERTISING RESEARCH

The integration of AI into advertising research has the potential to mark a significant shift in the industry's trajectory. AI, if properly utilized, can emerge as the definitive next step. This transformation isn't merely about data crunching; it's about the profound capability of AI to predict consumer behaviors, enhance creative processes, and provide real-time evaluations of campaigns. At the same time, as the various case studies in this handbook show, we are currently in the stage where the impact of AI is predominantly centered around enhancing efficiency and reducing operational times. AI's primary appeal in its current form is its ability to streamline processes, automate tasks, and analyze large datasets swiftly. In advertising

research, this translates to quicker summaries and analyses, more efficient interpretation of data, and designing and writing reports.

Despite these advantages, there's a crucial need for caution in AI deployment. One of the primary concerns is the accuracy and reliability of AI systems. AI algorithms are only as good as the data they are trained on, and biased or incomplete data can lead to inaccurate or even harmful outcomes. In advertising research, the risk of errors necessitates a strong framework of human oversight. AI should be viewed as a tool that augments human capabilities, not as a replacement. Human judgment is essential to interpret AI-generated insights correctly and make final decisions, especially in nuanced or complex scenarios that AI might not fully comprehend.

The promise of AI in advertising research, to offer deeper connections and more effective engagement with audiences, is tempered by these complexities. The responsibility falls on industry players to navigate these challenges carefully and to approach this promising but uncertain future with both optimism and a healthy dose of skepticism. Organizations venturing into AI in advertising research must first establish a foundational understanding of its relevance and applications, which is a complex and resource-intensive process. It requires substantial investment in infrastructure, training, and possibly collaboration with AI specialists. Initiating this integration with pilot projects might mitigate some risks, but it still presents a gradual and potentially arduous transition.

2. CONTINUED LEARNING AND ADAPTATION

In the rapidly evolving landscape of AI, static knowledge isn't enough. Continuous learning and adaptation are no longer optional but mandatory. As AI technologies mature and new methodologies emerge, there's an incessant need for professionals in advertising research to stay updated. Training doesn't conclude once an AI system is implemented; in fact, that's often just the beginning. Regular workshops, certifications, and seminars centered around AI advancements are crucial to ensure teams are equipped with the latest tools and knowledge.

Moreover, AI in advertising research is not a one-size-fits-all solution. As consumer behaviors evolve and industry trends shift, AI models and strategies need periodic re-evaluation and fine-tuning. Adaptation also means being open to feedback, both from within the organization and from the target audience. It's about acknowledging that, sometimes, human intuition can identify what raw data might miss. Such insights can then be fed back into AI systems, refining them further.

One of the promising avenues for this continued learning is the establishment of cross-collaborations with AI research institutions, tech startups, and other industries already leveraging AI. These collaborations can offer fresh perspectives, innovative solutions, and a broader understanding of AI's potential. They serve as a reminder that in the world of AI, learning is a collaborative and ever-evolving journey.

Furthermore, as AI-driven solutions become more mainstream, there's an inherent responsibility for organizations to educate not just their internal teams but also their audience. Transparency about how AI is used in research, the kind of data it processes, and its decision-making mechanisms can foster trust and bridge the knowledge gap. It paves the way for a

more informed and engaged consumer base, which in turn can provide valuable feedback, ensuring that AI tools are both effective and ethical.

In wrapping up, while the integration of AI into advertising research is transformative, its true power lies in an organization's commitment to perpetual learning and agile adaptation. Embracing this ethos ensures that AI doesn't just amplify current efforts, but continually reshapes and redefines the very contours of advertising research.

APPENDIX: GLOSSARY OF AI TERMS IN ADVERTISING RESEARCH

1. Algorithm: A set of rules or instructions given to an AI or machine learning model to perform a specific task. The IAB (2021) define it as follows: An algorithm is a sequence of well-defined computer instructions that solve a problem or perform a computation such as calculations, data processing, or automated reasoning.

2. Analytics: The discovery, interpretation, and communication of meaningful patterns in data.

3. Artificial Neural Network (ANN): A computing system made up of interconnected nodes that processes information similarly to how neurons in the human brain work. Commonly used in deep learning. IAB (2021) definition: Artificial neural networks (ANN) combine algorithms and computational power to process problems by mimicking biological neural networks' form and function like our brains.

4. **Backpropagation:** A method used in training neural networks by adjusting weights based on the error from the predicted output.

5. Bias: Pre-existing beliefs or distortions present in data or in the design of an algorithm which can lead to unfair or skewed results. In ML this refers to an algorithm's error due to erroneous assumptions in the learning process.

6. Big Data: Extremely large datasets that may be analyzed computationally to reveal patterns, trends, and associations.

7. Chatbot: A software application designed to simulate conversation with human users, especially over the internet.

8. Classification: A supervised learning task where the goal is to predict the categorical class labels of new instances.

9. Clustering: An unsupervised learning task where the aim is to group similar instances based on certain features.

10. Convolutional Neural Networks (CNNs): A class of deep neural networks, highly effective for processing data with a grid-like topology, such as images. CNNs utilize convolutional layers, pooling layers, and fully connected layers to automatically and adaptively learn spatial

hierarchies of features from input images. They excel in tasks like image and video recognition, image classification, medical image analysis, and natural language processing. The strength of CNNs lies in their ability to detect local patterns, such as edges in images, and their capacity to handle high-dimensional data, making them a cornerstone of computer vision and deep learning applications.

11. Data Mining: The process of discovering patterns and knowledge from large datasets.

12. Deep Learning: A subset of machine learning that uses neural networks with many layers.

13. Feature: An individual measurable property of the phenomenon being observed.

14. Genetic Algorithm: A search heuristic inspired by the process of natural selection.

15. **Generalization:** The model's ability to give accurate predictions for previously unseen data.

16. Generative Adversarial Network (GAN): A class of machine learning frameworks designed by Ian Goodfellow and his colleagues in 2014. GANs consist of two neural networks, a generator and a discriminator, which compete against each other. The generator creates data samples, while the discriminator evaluates them against real data. The goal of the generator is to produce data indistinguishable from real data, and the discriminator's goal is to correctly differentiate between the two. This adversarial process enhances the performance of both networks, enabling GANs to generate high-quality, realistic synthetic data, widely used in image, video, and voice generation.

17. Hyperautomation: Hyperautomation applies advanced technologies such as robotic process automation, AI, machine learning, and process mining to augment human workers and extend automation beyond traditional capabilities. Unlike simpler tools like macros or isolated scripts, hyperautomation tackles more complex, cognitive tasks, leading to more impactful automation processes.

18. Image Recognition: The ability of software to identify objects, places, people, or actions in images.

19. Keras is an open-source neural network library written in Python. It's designed to enable fast experimentation with deep neural networks and focuses on being user-friendly, modular, and extensible. Initially developed as an interface for TensorFlow, Keras now supports multiple backends, including TensorFlow, Microsoft Cognitive Toolkit (CNTK), and Theano. Keras simplifies the process of building and training deep learning models with its high-level building blocks for creating and training neural networks. It's widely used in both academia and industry and is known for its ease of use and flexibility, making it a popular choice for beginners and experts in machine learning and deep learning.

20. Knowledge Graph: A knowledge base that uses a graph-structured data model or topology to represent and link information.

21. Jupyter Notebooks: Jupyter Notebooks are an open-source web application that allows you to create and share documents containing live code, equations, visualizations, and narrative text. Widely used in data science, scientific computing, and machine learning, Jupyter Notebooks support various programming languages, including Python, R, and Julia.

They are particularly valued for their interactivity, ease of use, and ability to integrate code, explanatory text, and visual outputs in a single document. This makes them an excellent tool for exploratory data analysis, collaborative projects, educational purposes, and presenting computational workflows.

22. Machine Learning (ML): A subset of AI that allows computers to learn from data without being explicitly programmed.

23. Natural Language Processing (NLP): A field of AI that focuses on the interaction between computers and humans through natural language.

24. Overfitting: A modeling error that occurs when a machine learning model is tailored too closely to the training data and performs poorly on new, unseen data.

25. Perceptual Maps: Also known as positioning maps, are visual representations used in marketing to show how consumers perceive a product or brand in comparison to its competitors. These maps are typically based on consumer perceptions of certain attributes (like quality, price, or performance) and help marketers understand how their products are positioned in the market. By plotting products on a two-dimensional grid, perceptual maps illustrate the relationships between market competitors and provide insights into potential market gaps or areas for competitive advantage. This tool is invaluable for strategic planning in product development, branding, and advertising.

26. Predictive Analytics: Using historical data to predict future outcomes.

27. PyTorch: An open-source machine learning library developed by Facebook's AI Research lab. It's widely used for applications in computer vision, natural language processing, and deep learning. PyTorch is known for its flexibility and ease of use, particularly in the development and training of deep learning models. It provides dynamic computational graphing, allowing for intuitive and straightforward model building and modification. PyTorch also features strong GPU acceleration support and has a growing community, making it a popular choice among researchers and developers for both experimentation and production in AI and deep learning projects.

28. Regression: A type of supervised learning where the aim is to predict continuous values.

29. Reinforcement Learning: A type of machine learning where agents learn by interacting with their environment and receiving feedback in the form of rewards or penalties.

30. Sentiment Analysis: Using NLP to determine whether a piece of text is positive, negative, or neutral in tone.

31. Supervised Learning: A type of machine learning where the model is trained using labeled data.

32. Support Vector Machine (SVM): A powerful and versatile supervised machine learning algorithm, primarily used for classification and regression challenges. It works by finding the best hyperplane that separates data points into different classes in the feature space. The strength of SVM lies in its ability to handle linear and non-linear data using kernel functions. SVMs are popular in applications like image classification, text categorization,

and bioinformatics, due to their effectiveness in handling high-dimensional data and their robustness against overfitting, especially in cases where the number of dimensions exceeds the number of samples.

33. TensorFlow: An open-source software library for numerical computation using data flow graphs. Developed by the Google Brain team, it's widely used in machine learning and deep learning for building and training neural networks. TensorFlow's flexible architecture allows for easy deployment of computation across various platforms (CPUs, GPUs, and TPUs), facilitating both research and production use. It supports a wide range of tasks, primarily focused on training and inference of deep neural networks in areas like computer vision, natural language processing, and predictive analytics. TensorFlow's user-friendly interface and extensive community support have made it a popular tool in the AI field.

34. Text-to-Speech/Speech-to-Text: Text-to-speech and speech-to-text technologies enable the accurate recognition and translation of oral language into textual form, and the effective transformation of written text into audible speech.

35. Time Series Forecasting: Using machine learning to predict future values based on previously observed values.

36. Training Data: The data on which a machine learning model is trained. Test dataset is used to assess the performance of the model.

37. Turing Test: The Turing Test, proposed by Alan Turing in 1950, is a method for determining whether a machine exhibits intelligent behavior equivalent to, or indistinguishable from, that of a human. In this test, a human judge engages in a natural language conversation with one human and one machine, both of which try to appear human. If the judge cannot reliably tell the machine from the human, the machine is said to have passed the test. The Turing Test has been a fundamental concept in the philosophy of artificial intelligence, prompting debates about the nature of intelligence and the potential of machines to exhibit human-like consciousness.

38. Unsupervised Learning: A type of machine learning where the model learns from data without labeled responses.

39. Voice Recognition: The ability of a machine or program to receive and interpret dictation or to understand and carry out spoken commands.

40. Weights: Values in a neural network that transform input data within the network's layers.

REFERENCES

Abdelgaber, N., & Nikolopoulos, C. (2020). Overview on quantum computing and its applications in artificial intelligence. *2020 IEEE Third International Conference on Artificial Intelligence and Knowledge Engineering (AIKE)*, 198-199.

Adi, E., Anwar, A., Baig, Z., & Zeadally, S. (2020). Machine learning and data analytics for the IoT. *Neural Computing and Applications*, *32*, 16205-16233.

Agbehadji, I. E., Awuzie, B. O., Ngowi, A. B., & Millham, R. C. (2020). Review of big data analytics, artificial intelligence and nature-inspired computing models towards accurate detection of COVID-19 pandemic cases and contact tracing. *International Journal of Environmental Research and Public Health*, *17*(15), 5330.

Aggarwal, P., & McGill, A. L. (2007). Is that car smiling at me? Schema congruity as a basis for evaluating anthropomorphized products. *Journal of Consumer Research*, *34*(4), 468-479.

Agrawal R., Gehrke J., Gunopulos D., & Raghavan P., (1994). Fast algorithms for mining association rules. *Proceedings of the International Joint Conference on Very Large Data Bases, Santiago Chile,* 1215, 487–499.

Agrawal, R., Imieliński, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, 207-216).

Agrawal, D., & Schorling, C. (1996). Market share forecasting: An empirical comparison of artificial neural networks and multinomial logit model. Journal of Retailing, 72(4), 383–408.

Ali, M.J. & Djalilian, A. (2023). Readership awareness series–paper 4: Chatbots and ChatGPT-ethical considerations in scientific publications. *Seminars in Ophthalmology*, 38(5), 403-404.

Ameen, N., Tarhini, A., Shah, M. H., Madichie, N., Paul, J., & Choudrie, J. (2020). Keeping customers' data secure: A cross-cultural study of cybersecurity compliance among the GenMobile workforce. *Computers in Human Behavior*, 114, 106531.

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1985). Spin-glass models of neural networks. *Physical Review A*, *32*(2), 1007.

Analytic Partners. (2023, Q3). Generative AI in marketing measurement & optimization. Analytic Partners.

Ansari, A., & Riasi, A. (2016). Modelling and evaluating customer loyalty using neural networks: Evidence from startup insurance companies. *Future Business Journal*, 2(1), 15–30.

Anzai, Y. (2012). Pattern recognition and machine learning. Academic Press, Inc.

Araujo, T., Copulsky, J. R., Hayes, J. L., Kim, S. J., & Srivastava, J. (2020). From Purchasing Exposure to Fostering Engagement: Brand-Consumer Experiences in the Emerging Computational Advertising Landscape. *Journal of Advertising*, *49*(4), 428–445.

ARF. (2023, March). *ARF immersive advertising study.* ARF.

ARF. (2023, July). *ARF Attention Measurement Validation Initiative: Literature review*. ARF.

ARF. (2023, September). *ARF Attention Measurement Validation Initiative: Phase 1*. ARF.

Arora, A. (2023, July). Moravec's paradox and the fear of job automation in health care. The Lancet.

Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., ... & Zhang, Z. J. (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, *19*, 305-321.

Athaluri, S. A., Manthena, S. V., Kesapragada, V. K. M., Yarlagadda, V., Dave, T., & Duddumpudi, R. T. S. (2023). Exploring the boundaries of reality: investigating the phenomenon of artificial intelligence hallucination in scientific writing through ChatGPT references. *Cureus*, *15*(4), e37432.

Bakpayev, M., Baek, T. H., van Esch, P., & Yoon, S. (2020). Programmatic creative: AI can think but it cannot feel. *Australasian Marketing Journal, 30*(1), 90–95.

Ballestar, M. T., Grau-Carles, P., & Sainz, J. (2019). Predicting customer quality in e- commerce social networks: A machine learning approach. *Review of Managerial Science*, 13(3), 589–603.

Baltrušaitis, T., Robinson, P., & Morency, L. P. (2016). Openface: An open source facial behavior analysis toolkit. *2016 IEEE winter conference on applications of computer vision (WACV)*, 1-10.

Bandgar, S. (2021, Oct). Support Vector Machine. Medium.

Baumann, C., Elliott, G., & Burton, S. (2012). Modeling customer satisfaction and loyalty: Survey data versus data mining. *Journal of Services Marketing*, 26(3), 148–157.

Bayrak Meydanoğlu, E. S., Çilingirtürk, A. M., Öztürk, R., & Klein, M. (2020). An Empirical Cross-Country Study On Consumers' Attitude Towards Augmented Reality Advertising. *Business and Management Studies: An International Journal*, 8(2), 1424-1454.

Bejou, D., Wray, B., & Ingram, T. N. (1996). Determinants of relationship quality: An artificial neural network analysis. *Journal of Business Research*, 36(2), 137–143.

Belanche D., Casaló, L.V. & Flavián, C. (2019). Artificial Intelligence in Fintech: Understanding Robo-Advisors Adoption Among Customers. *Industrial Management and Data Systems*, 119(7), 1411-1430.

Belk, R. (2016). Understanding the robot: Comments on Goudey and Bonnin (2016). *Recherche et Applications en Marketing (English Edition)*, 31(4), 83-90.

Bock, D. E., Wolter, J. S., & Ferrell, O. C. (2020). Artificial intelligence: disrupting what we know about services. *Journal of Services Marketing*, 34(3), 317-334.

Boddington, P. (2017). *Towards a code of ethics for artificial intelligence*. Springer International Publishing.

Boerman, S. C., & Smit, E. G., (2023) Advertising and privacy: an overview of past research and a research agenda, *International Journal of Advertising*, 42(1), 60-68.

Bohr, A., & Memarzadeh, K. (2020) The rise of artificial intelligence in healthcare applications. In *Artificial intelligence in healthcare Academic Press*, 25-60.

Bogue, R. (2014). The role of artificial intelligence in robotics. *Industrial Robot: An International Journal*, *41*(2), 119-123.

Bostrom, N., & Yudkowsky, E. (2018). The ethics of artificial intelligence. In *Artificial intelligence safety and security* (pp. 57-69). Chapman and Hall/CRC.

Bova, F., Goldfarb, A., & Melko, R. G. (2021). Commercial applications of quantum computing. *EPJ Quantum Technology*, *8*(1), 2.

Bragazzi, N. L., Dai, H., Damiani, G., Behzadifar, M., Martini, M., & Wu, J. (2020). How big data and artificial intelligence can help better manage the COVID-19 pandemic. *International Journal of Environmental Research and Public Health*, *17*(9), 3176.

Brand, J., Israeli, A., & Ngwe, D. (2023). Using GPT for market research. *Available at SSRN 4395751*.

Briggs, C., & Briggs, R. (2024). The AI Conundrum: Harnessing the Power of AI for Your Organization Profitably and Safely. MIT Press.

Brock, D. C. (2018). Learning from Artificial Intelligence's Previous Awakenings: The History of Expert Systems. *AI Magazine*, *39*(3), 3-15.

Brossard, M., Gatto, G., Gentile, A., Merle, T., and Wlezien, C. (2020, Feb). How generative design could reshape the future of product development. *McKinsey & Company*.

Buchanan, B. G. (2005). A (very) brief history of artificial intelligence. *AI Magazine*, *26*(4), 53-53.

Butterworth, M. (2018). The ICO and artificial intelligence: The role of fairness in the GDPR framework. *Computer Law & Security Review*, 34(2), 257-268.

Calvano, E., Calzolari, G., Denicolò, V. & Pastorello, S. (2019). Algorithmic pricing what implications for competition policy? *Review of Industrial Organization*, 55(1), 155-171.

Campbell, M. (2023, March). Integrating AI into market research. ResearchWorld.

Campbell, C., Plangger, K., Sands, S., Kietzmann, J., & Bates, K. (2022). How deepfakes and artificial intelligence could reshape the advertising industry: The coming reality of AI fakes and their potential impact on consumer behavior. *Journal of Advertising Research*, *62*(3), 241-251.

Cao, L. (2017). Data science: A comprehensive overview. ACM Computing Surveys (CSUR), 50(3), 43.

Carbonneau, R., Laframboise, K., & Vahidov, R. (2008). Application of machine learning techniques for supply chain demand forecasting. *European Journal of Operational Research*, 184, 1140–1154.

Casabayó, M., Agell, N., Aguado, J. C. (2004). Using AI Techniques in the Grocery Industry: Identifying the Customers Most Likely to Defect. *International Review of Retail, Distribution and Consumer Research*, 14(3), 295-308.

Cascio, R., Mariadoss, B. J., & Mouri, N. (2010). The impact of management commitment alignment on salespersons' adoption of sales force automation technologies: An empirical investigation. *Industrial Marketing Management*, 39(7), 1088–1096.

Chan, S. L., & Ip, W. H. (2011). A Dynamic Decision Support System to Predict the Value of Customer for New Product Development. *Decision Support Systems*, 52(1), 178-188.

Chen, G., Xie, P., Dong, J., & Wang, T. (2019). Understanding Programmatic Creative: The Role of AI. *Journal of Advertising, 48*(4), 347–355.

Cherviakova, V., & Cherviakova, T. (2018). Value Opportunities for Automotive Manufacturers in Conditions of Digital Transformation of the Automotive Industry. *Journal of Applied Economic Sciences*, 13(8), 2351-2362.

Chica, M., Cordón, Ó., Damas, S., Iglesias, V., & Mingot, J. (2016). Identimod: Modeling and managing brand value using soft computing. *Decision Support Systems*, 89, 41-55.

Chiu, C. C., Sainath, T. N., Wu, Y., Prabhavalkar, R., Nguyen, P., Chen, Z., Kannan, A., Weiss, R. J., Rao, K., & Goninam, E. (2018). State-of-the-art speech recognition with sequence-to-sequence models. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4774–4778.

Choi, E. P. H., Lee, J. J., Ho, M. H., Kwok, J. Y. Y., & Lok, K. Y. W. (2023). Chatting or cheating? The impacts of ChatGPT and other artificial intelligence language models on nurse education. *Nurse Education Today*, *125*, 105796-105796.

Choi, Y. S., & Yu, S. Y. (2019). A Study on the Causal Relationship Between Attributes of VR Ad and Advertising Effectiveness: Mediating Effect of the Consumer's Innovation. *Journal of Digital Convergence*, *17*(2), 37-45.

Chong, A. Y. L., Li, B., Ngai, E. W., Ch'ng, E., & Lee, F. (2016). Predicting online product sales via online reviews, sentiments, and promotion strategies. *International Journal of Operations & Production Management*, 36(4), 358–383.

Chui, M., Yee, L., Hall, B., Singla, A., & Sukharevsky, A. (2023, August 1). *The state of AI in 2023: Generative AI's breakout year.* McKinsey.

Coffee, P. (2023, December). An Anticipated Wave of AI Specialist Jobs Has Yet to Arrive. The Wall Street Journal.

Coffin, J. (2022). Asking Questions of AI Advertising: A Maieutic Approach. *Journal of Advertising, 51*(5), 608–623.

Coldewey, D., & Lardinois, F. (2023, Feb). AI is eating itself: Bing's AI quotes COVID disinfo sourced from ChatGPT. TechCrunch+.

Cooke, M., & Passingham, J. (2022.) Do androids dream of electric research? WARC.

Cunningham, P., Cord, M., & Delany, S. J. (2008). Supervised learning. In *Machine learning techniques for multimedia: case studies on organization and retrieval* (pp. 21-49). *Berlin, Heidelberg: Springer Berlin Heidelberg.*

De Cicco, R., Silva, S. & Alparone, F. (2020). Millennials' attitude toward chatbots: An experimental study in a social relationship perspective. *International Journal of Retail and Distribution Management*, available online ahead of print.

DeJong, G., & Mooney, R. (1986). Explanation-based learning: An alternative view. *Machine learning*, *1*, 145-176.

Deng, L., & Liu, Y. (Eds.). (2018). *Deep learning in natural language processing*. Berlin: Springer.

Deng, S., Tan, C.-W., Wang, W., & Pan, Y. (2019). Smart Generation System of Personalized Advertising Copy and Its Application to Advertising Practice and Research. *Journal of Advertising*, *48*(4), 356–365.

Devagiri, J. S., Paheding, S., Niyaz, Q., Yang, X., & Smith, S. (2022). Augmented Reality and Artificial Intelligence in industry: Trends, tools, and future challenges. *Expert Systems with Applications*, 118002.

Deveau, R., Griffin, S. J., & Reis, S. (2023, May). AI-powered marketing and sales reach new heights with generative AI. McKinsey.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

D'Haen, J., & Van den Poel, D. (2013). Model-supported business-to-business prospect prediction based on an iterative customer acquisition framework. *Industrial Marketing Management*, 42(4), 544-551.

Diaz, J. (2017). How marketers are using machine learning to cross-sell and upsell. Principa

Dignum, V. (2018). Ethics in artificial intelligence: introduction to the special issue. *Ethics Inf Technol* 20, 1–3.

Donthu, N., & Gustafsson, A. (2020). Effects of COVID-19 on business and research. *Journal of Business Research*, 117, 284-289.

Du, X., Jiao, J., & Tseng, M. M. (2003). Identifying customer need patterns for customization and personalization. *Integrated Manufacturing Systems*, 14(5), 387-396.

European Commission. (2021). *EU-US Trade and Technology Council Statement 2021*. EC.

Erdem, E., & Sidlova, V. (2023, November). The future of generative AI in advertising: Efficiency without effectiveness? Kantar.

Erevelles, S., Fukawa, N., & Swayne, L. (2016). Big Data consumer analytics and the transformation of marketing. *Journal of Business Research*, 69(2), 897–904.

Etzioni, A., & Etzioni, O. (2017). Incorporating ethics into artificial intelligence. *The Journal of Ethics*, 21(4), 403-418.

Farrokhnia, M., Banihashem, S. K., Noroozi, O., & Wals, A. (2023). A SWOT analysis of ChatGPT: Implications for educational practice and research. *Innovations in Education and Teaching International*, 1-15.

Fatima, M., & Pasha, M. (2017). Survey of machine learning algorithms for disease diagnostic. *Journal of Intelligent Learning Systems and Applications*, *9*(01), 1-16.

Firican, G. The history of big data. *Lights On Data*.

Fish, K. E., Barnes, J. H., & Aiken, M. W. (1995). Artificial neural networks: a new methodology for industrial market segmentation. *Industrial Marketing Management*, 24(5), 431-438.

Fish, K. E., Johnson, J. D., Dorsey, R. E., & Blodgett, J. G. (2004). Using an artificial neural network trained with a genetic algorithm to model brand share. *Journal of Business Research*, 57(1), 79-85.

Flaherty, K. E., Lassk, F., Lee, N., Marshall, G. W., Moncrief, W. C., Mulki, J. P., & Pullins, E. B. (2018). Sales scholarship: honoring the past and defining the future (Key takeaways from the 2018 American Marketing Association Faculty Consortium: New Horizons in Selling and Sales Management). *Journal of Personal Selling & Sales Management*, 38(4), 413-421.

Floridi, L. & Cowls, J. (2019). A Unified Framework of Five Principles for AI in Society. *Harvard Data Science Review,* 1(1).

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). *AI4People — An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles and Recommendations.* (AI4People).

Fong, H., Kumar, V., & Sudhir, K. (2021). A theory-based interpretable deep learning architecture for music emotion. *Available at SSRN 4025386.*

Ford, J., Jain, V., Wadhwani, K., & Gupta, D. G. (2023). AI advertising: An overview and guidelines. *Journal of Business Research*, *166*, 114124.

Fried, I. (2023 Feb). ChatGPT sparks surge of AI detection tools. Axios.

Fujiyoshi, H., Hirakawa, T., & Yamashita, T. (2019). Deep learning-based image recognition for autonomous driving. *IATSS Research*, 43(4), 244–252.

Frąckiewicz, M. (2023). The Evolution of AI: Tracing the History and Progress of Decision Trees in Artificial Intelligence. Artificial Intelligence, TS2 Space.

Gardner, E. (1988). The space of interactions in neural network models. *Journal of Physics A: Mathematical and General*, 21(1), 257.

Gardner, E. (1987). Maximum storage capacity in neural networks. Europhysics Letters, 4(4), 481.

Géron, A. (2022). Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow. O'Reilly Media, Inc.

Ghahramani, Z. (2003). Unsupervised learning. In *Summer school on machine learning* (pp. 72-112). Berlin, Heidelberg: Springer Berlin Heidelberg.

Girotra, K., Meincke, L., Terwiesch, C., & Ulrich, K. T. (2023). Ideas are dimes a dozen: Large language models for idea generation in innovation. *Available at SSRN 4526071.*

Goebel, R., Chander, A., Holzinger, K., Lecue, F., Akata, Z., Stumpf, S., ... & Holzinger, A. (2018). Explainable AI: the new 42?. Goebel, R. *et al.* (2018). Explainable AI: The New 42?. In: Holzinger, A., Kieseberg, P., Tjoa, A., Weippl, E. (eds) Machine Learning and Knowledge Extraction. CD-MAKE 2018. *Lecture Notes in Computer Science()*, vol 11015. Springer, Cham. https://doi.org/10.1007/978-3-319-99740-7_21

Goertzel, B. (2014). Artificial general intelligence: concept, state of the art, and future prospects. *Journal of Artificial General Intelligence*, *5*(1), 1.

Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, *32*, 100577.

Goldstein, I., & Papert, S. (1977). Artificial Intelligence, Language and the Study of Knowledge. *Cognitive Science*, 1(1).

Gong, J., Abhishek, V., & Li, B. (2018). Examining the Impact of Keyword Ambiguity on Search Advertising Performance. *MIS Quarterly*, *42*(3), 805-A14.

Gozalo-Brizuela, R., & Garrido-Merchan, E. C. (2023). ChatGPT is not all you need. A State of the Art Review of large Generative AI models. *arXiv preprint arXiv:2301.04655.*

Gralpois, B. (2023, March). The future of AI in advertising and client/agency relationships. ANA.

Grewal, D., Hulland, J., Kopalle, P. K., & Karahanna, E. (2020). The future of technology and marketing: a multidisciplinary perspective. *Journal of the Academy of Marketing Science*, 48, 1-8.

Griffith, D., Kiessling, T. & Dabić, M. (2012). Aligning strategic orientation with local market conditions: Implications for subsidiary knowledge management. *International Marketing Review*, 29(4), 379-402.

Grudin, J., & Jacques, R. (2019). Chatbots, humbots, and the quest for artificial general intelligence. *Proceedings of the 2019 CHI conference on human factors in computing systems*, 1-11.

Guitart, I. A., Hervet, G., & Gelper, S. (2020). Competitive advertising strategies for programmatic television. *J. of the Acad. Mark. Sci.*, *48*(4), 753–775.

Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G. Z. (2019). XAI—Explainable artificial intelligence. *Science Robotics*, *4*(37), eaay7120.

Ha, K., Cho, S., & MacLachlan, D. (2005). Response models based on bagging neural networks. *Journal of Interactive Marketing*, 19(1), 17-30.

Haleem, A., Javaid, M., Qadri, M. A., Singh, R. P., & Suman, R. (2022). Artificial intelligence (AI) applications for marketing: A literature-based study. *International Journal of Intelligent Networks*, 3, 119-132.

Hamid, S. A., & Iqbal, Z. (2004). Using neural networks for forecasting volatility of S&P 500 Index futures prices. *Journal of Business Research*, 57(10), 1116–1125.

Han, J., Pei, J., & Tong, H. (2022). *Data mining: concepts and techniques*. Morgan kaufmann.

Han, J., Pei, J., & Yin, Y. (2000). Mining frequent patterns without candidate generation. *ACM SIGMOND Record*, *29*(2), 1-12.

Hardcastle, K. (2023, November). How generative AI is redefining the game for marketers, and what can they do about it? WARC.

Haryanto, J. O., Silva, M., & Moutinho, L. (2015). Neural network approach to understanding the children's market. *European Journal of Marketing*, 49(3–4), 372–397.

Hassani, H., & Silva, E. S. (2023). The role of ChatGPT in data science: how ai-assisted conversational interfaces are revolutionizing the field. *Big Data and Cognitive Computing*, *7*(2), 62.

Hayes, J. L., Britt, B. C., Evans, W., Rush, S. W., Towery, N. A., & Adamson, A. C. (2021). Can social media listening platforms' artificial intelligence be trusted? Examining the accuracy of Crimson Hexagon's (now Brandwatch Consumer Research's) AI-Driven analyses. *Journal of Advertising*, *50*(1), 81-91.

Henke, N.; Bughin, J.; Chui, M.; Manyika, J.; Saleh, T.;Wiseman, B.; Sethupathy, G. (2016, Dec). The Age of Analytics: Competing in a Data-Driven World. McKinsey Global Institute.

Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural computation*, *18*(7), 1527-1554.

Ho, C., Mu, J., Timpone, R., & Souchal, C. (2022, July). Beyond the hype- Innovation predictions in the era of machine learning. Ipsos.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554-2558.

Hsu, T., & Lu, Y. (2023, July). A blessing and a boogeyman: Advertisers warily embrace A.I. *The New York Times*.

Huang, MH., & Rust, R.T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science* , 49, 30–50.

Huang, M. H., & Rust, R. T. (2020). Engaged to a Robot? The Role of AI in Service. *Journal of Service Research*, 24(1).

Huang, M. H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155-172.

Hughes-Castleberry, K. (2023, January). Inside Quantum Technology's Inside Scoop: Quantum and Advertising. Inside Quantum Technology News.

Huh, J., Kim, H., Rath, B., Lu, X., & Srivastava, J. (2020). You reap where you sow: A trust-based approach to initial seeding for viral advertising. *International Journal of Advertising*, *39*(7), 963–989.

Huh, J., Nelson, M. R., & Russell, C. A. (2023). ChatGPT, AI advertising, and advertising research and education. *Journal of Advertising*, *52*(4), 477-482.

Huff, E., & Bonde, A. (2022, May). AI meets consumer insights: Welcome to the era of AICI. Ipsos.

IAB AI Standards Working Group. (2021, March). Artificial intelligence use cases and best practices for marketing. IAB.

Ipsos (2023, July). What the Future. Intelligence.

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. Journal of Artificial Intelligence Research, 4, 237-285.

Khan, R. A., Jawaid, M., Khan, A. R., & Sajjad, M. (2023). ChatGPT-Reshaping medical education and clinical management. *Pakistan Journal of Medical Sciences*, *39*(2), 605 - 607.

Kim, Y., Street, W. N., Russell, G. J., & Menczer, F. (2005). Customer targeting: A neural network approach guided by genetic algorithms. *Management Science*, 51(2), 264–276.

Kitsopoulou, M., & Lappas, G. (2023, November). AR/VR technologies in advertising: A scoping review on empirical studies on the effectiveness of AR/VR in advertising. In *AIP Conference Proceedings* (Vol. 2909, No. 1). AIP Publishing.

Kerwin, A. M. (2023, Oct). Spotlight US: How US marketers are starting to use Gen AI. WARC.

Kibria, M. G., Nguyen, K., Villardi, G. P., Zhao, O., Ishizu, K., & Kojima, F. (2018). Big data analytics, machine learning, and artificial intelligence in next-generation wireless networks. *IEEE access*, *6*, 32328-32338.

Kietzmann, J., Paschen, J., & Treen, E. (2018). Artificial intelligence in advertising: How marketers can leverage artificial intelligence along the consumer journey. *Journal of Advertising Research*, 58(3), 263-267.

Kim, Y., & Street, W. N. (2004). An Intelligent System for Customer Targeting: A Data Mining Approach. *Decision Support Systems*, 37(2), 215-228.

Kitchens, B., Dobolyi, D., Li, J., & Abbasi, A. (2018). Advanced customer analytics: Strategic value through integration of relationship-oriented big data. *Journal of Management Information Systems*, 35(2), 540–574.

Kodinariya, T. M., & Makwana, P. R. (2013). Review on determining number of Cluster in K-Means Clustering. *International Journal*, 1(6), 90-95.

Kolbjørnsrud, V., Amico, R., & Thomas, R. J. (2017). Partnering with AI: how organizations can win over skeptical managers. *Strategy & Leadership,* 45(1), 37-43.

Kononova, A., Kim, W., Joo, E., & Lynch, K. (2020). Click, click, ad: The proportion of relevant (vs. irrelevant) ads matters when advertising within paginated online content. *International Journal of Advertising, 39*(7), 1031–1058.

Korganbekova, M., and C. Zuber (2023). Balancing User Privacy and Personalization. *Work In Progress*.

Korteling, J. H., van de Boer-Visschedijk, G. C., Blankendaal, R. A., Boonekamp, R. C., & Eikelboom, A. R. (2021). Human-versus artificial intelligence. *Frontiers in Artificial Intelligence*, *4*, 622364.

Kotsiantis, S. B., Zaharakis, I. D., & Pintelas, P. E. (2006). Machine learning: a review of classification and combining techniques. *Artificial Intelligence Review*, 26, 159-190.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems, 25*.

Kuligowski, K. (2023, October). Facial Recognition Advertising: The New Way to Target Ads to Consumers. Business News Daily.

Kumar, Y., Koul, A., Singla, R., & Ijaz, M. F. (2022). Artificial intelligence in disease diagnosis: a systematic literature review, synthesizing framework and future research agenda. *Journal of Ambient Intelligence and Humanized Computing*, 1-28.

Kumar, A., Paul, J., & Unnithan, A. B. (2020). Masstige marketing: A review, synthesis and research agenda. *Journal of Business Research*, 113, 384–398.

Kushwaha, S., Bahl, S., Bagha, A. K., Parmar, K. S., Javaid, M., Haleem, A., & Singh, R. P. (2020). Significant applications of machine learning for COVID-19 pandemic. *Journal of Industrial Integration and Management*, *5*(04), 453-479.

Lau, H. C., Nakandala, D., Zhao, L., & Lai, I. K. (2015). Using fuzzy logic approach in estimating individual guest loyalty level for international tourist hotels. *International Journal of Services Technology and Management*, 21(1–3), 127–145.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436-444.

Lee, H., & Cho, C.-H. (2020). Uses and gratifications of smart speakers: Modelling the effectiveness of smart speaker advertising. *International Journal of Advertising, 39*(7), 1150–1171.

Lei, N., & Moon, S. K. (2015). A Decision Support System for market-driven product positioning and design. Decision Support Systems, 69, 82-91.

Leminen, S., Rajahonka, M., Westerlund, M., & Wendelin, R. (2018). The future of the Internet of Things: Toward heterarchical ecosystems and service business models. *Journal of Business & Industrial Marketing*, 33(6), 749–767.

Lewinski, P., Den Uyl, T. M., & Butler, C. (2014). Automated facial coding: Validation of basic emotions and FACS AUs in FaceReader. *Journal of Neuroscience, Psychology, and Economics*, *7*(4), 227.

Li, H. (2019). Special Section Introduction: Artificial Intelligence and Advertising. *Journal of Advertising*, *48*(4), 333–337.

Li, S. (2000). The development of a hybrid intelligent system for developing marketing strategy. Decision Support Systems, 27(4), 395-409.

Li, K., & Du, T. C. (2012). Building a targeted mobile advertising system for location-based services. *Decision Support Systems*, *54*(1), 1–8.

Li, Y., Hou, M., Liu, H., & Liu, Y. (2012). Towards a theoretical framework of strategic decision, supporting capability and information sharing under the context of Internet of Things. *Information Technology and Management*, 13(4), 205–216.

Li, H., & Kannan, P. K. (2014). Attributing conversions in a multichannel online marketing environment: An empirical model and a field experiment. *Journal of Marketing Research*, *51*(1), 40-56.

Liang, W., Zhang, Y., Cao, H., Wang, B., Ding, D., Yang, X., ... & Zou, J. (2023). Can large language models provide useful feedback on research papers? A large-scale empirical analysis. *arXiv preprint arXiv:2310.01783.*

Lin, C., & Kunnathur, A. (2019). Strategic orientations, developmental culture, and big data capability. Journal of Business Research, 105, 49–60.

Lina, L. F., & Setiyanto, A. (2021). Privacy concerns in personalized advertising effectiveness on social media. *Sriwijaya International Journal of Dynamic Economics and Business*, *5*(2), 147-156.

Lindsay, R. K., Feigenbaum, E. A., Buchanan, B. G., & Lederberg, J. (1980). Applications of artificial intelligence for chemical inference: The DENDRAL project. New York: McGraw-Hill.

Lim, W. M., Gunasekara, A., Pallant, J. L., Pallant, J. I., & Pechenkina, E. (2023). Generative AI and the future of education: Ragnarök or reformation? A paradoxical perspective from management educators. *The International Journal of Management Education*, *21*(2), 100790.

Lippman, A. (1980). Movie-maps: An application of the optical videodisc to computer graphics. *Acm Siggraph Computer Graphics*, *14*(3), 32-42.

Liu, B., Hsu, W., & Ma, Y. (1998, August). Integrating classification and association rule mining. *Proceedings of the fourth international conference on knowledge discovery and data mining* , 80-86.

Liu-Thompkins, Y., Maslowska, E., Ren, Y., & Kim, H. (2020). Creating, metavoicing, and propagating: A road map for understanding user roles in computational advertising. *Journal of Advertising, 49*(4), 394–410.

Lo, F.-Y., & Campos, N. (2018). Blending internet-of-things (IoT) solutions into relationship marketing strategies. *Technological Forecasting and Social Change*, 137, 10–18.

López, G., Quesada, L., & Guerrero, L. A. (2018). Alexa vs. Siri vs. Cortana vs. Google Assistant: a comparison of speech-based natural user interfaces. *Advances in Human Factors and Systems Interaction: Proceedings of the AHFE 2017 International Conference on Human Factors and Systems Interaction,* 241-250.

Mahdavinejad, M.S., Rezvan, M., Barekatain, M., Adibi, P., Barnaghi, P., & Sheth A.P. (2018). Machine learning for internet of things data analysis: A survey. *Digital Communications and Networks, 4*(3), 161–75.

Makridakis, S. (2017). The forthcoming artificial intelligence (AI) revolution: Its impact on society and firms. *Futures, 90*, 46-60.

Malthouse, E. C., Hessary, Y. K., Vakeel, K. A., Burke, R., & Fudurić, M. (2019). An algorithm for allocating sponsored recommendations and content: Unifying programmatic advertising and recommender systems. *Journal of Advertising, 48*(4), 366–379.

Malthouse, E. C., & Li, H. (2017). Opportunities for and pitfalls of using big data in advertising research. *Journal of Advertising, 46*(2), 227-235.

Manole, L. (2023, January 13). *Top AI trends affecting the market research industry*. Greenbook.

Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., & Byers, A. H. (2011). *Big data: The next frontier for innovation, competition, and productivity.* McKinsey.

Marchand, A., & Marx, P. (2020). Automated product recommendations with preference-based explanations. *Journal of Retailing*, *96*(3), 328-343.

Mariani, M. M., Perez-Vega, R., & Wirtz, J. (2022). AI in marketing, consumer research and psychology: A systematic literature review and research agenda. *Psychology & Marketing*, *39*(4), 755-776.

Mateas, M. (2002). *Interactive drama, art and artificial intelligence*. Carnegie Mellon University.

Matz, S. C., Segalin, C., Stillwell, D., Müller, S. R., & Bos, M. W. (2019). Predicting the personal appeal of marketing images using computational methods. *Journal of Consumer Psychology, 29*(3), 370–390.

McCarthy, J., Minsky, M.L., Rochester, N., & Shannon, C.E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Magazine, 27*(4), 12.

Meiseberg, B. (2016). The effectiveness of e-tailers' communication practices in stimulating sales of niche versus popular products. *Journal of Retailing, 92*(3), 319-332.

Metz, C. (2023, November 16). Chatbots may 'hallucinate' more often than many realize. *The New York Times*.

Miralles-Pechuán, L., Ponce, H., & Martínez-Villaseñor, L. (2018). A novel methodology for optimizing display advertising campaigns using genetic algorithms. *Electronic Commerce Research and Applications, 27*, 39–51.

Mittal, N., Saif, I., & Ammanath, B. (2022, October). *Fueling AI transformation: Four key actions powering widespread value from AI, right now. Deloitte's State of AI in the Enterprise, 5th edition report*. Deloitte.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing Atari with deep reinforcement learning. *arXiv:1312.5602*.

Moriuchi, E. (2019). Okay, Google! An empirical study on voice assistants on consumer engagement and loyalty. *Psychology & Marketing, 35*(5), 489-501.

Morris, R. J. T., & Truskowski, B. J. (2003). The evolution of storage systems. *IBM Systems Journal, 42*(2), 205-217.

Murtagh, F., & Contreras, P. (2012). Algorithms for hierarchical clustering: An overview. *WIREs Data Mining and Knowledge Discovery, 2*(1), 86-97.

Mustak, M., Salminen, J., Plé, L., & Wirtz, J. (2021). Artificial intelligence in marketing: Topic modeling, scientometric analysis, and research agenda. *Journal of Business Research, 124*, 389–404.

Naimark, M. (2006). Aspen the verb: Musings on heritage and virtuality. *Presence Journal*, special issue on Virtual Heritage.

Netzer, O., Feldman, R., Goldenberg, J., & Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science, 31*(3), 521–543.

Neumann, N., Tucker, C. E., & Whitfield, T. (2019). Frontiers: How effective is third-party consumer profiling? Evidence from field studies. *Marketing Science, 38*(6), 918–926.

Nilashi, M., bin Ibrahim, O., Ahmadi, H., & Shahmoradi, L. (2017). An analytical method for diseases prediction using machine learning techniques. *Computers & Chemical Engineering, 106*, 212-223.

Noble, W. S. (2006). What is a support vector machine? *Nature Biotechnology, 24*(12), 1565-1567.

Ostler, J., & Kalidas, A. (2023, March 31). *What large language models could mean for market research.* Kantar.

Ostler, J., Munoz, L. V., & Krithivasan, D. (2023, May 25). *What could generative AI mean for advertising and concept development?* Kantar.

Otter, D. W., Medina, J. R., & Kalita, J. K. (2020). A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, *32*(2), 604-624.

Palos-Sanchez, P., Saura, J. R., & Martin-Velicia, F. (2019). A study of the effects of programmatic advertising on users' concerns about privacy overtime. *Journal of Business Research*, *96*, 61–72.

Pantano, E., Giglio, S., & Dennis, C. (2019). Making sense of consumers' tweets: Sentiment outcomes for fast fashion retailers through Big Data analytics. *International Journal of Retail & Distribution Management*, *47*(9), 915–927.

Parachuk, T. (2021, October 27). *How voice-activated ads have captured the market.* Voices.

Paschen, J. (2020). Investigating the emotional appeal of fake news using artificial intelligence and human contributions. *Journal of Product and Brand Management*, *29*(2), 223-233.

Paschen, J., Wilson, M., & Ferreira, J. J. (2020). Collaborative intelligence: How human and artificial intelligence create value along the B2B sales funnel. *Business Horizons,* *63*(3), 403-414.

Paul, J. (2019). Masstige model and measure for brand management. *European Management Journal*, *37*(3), 299-312.

Pei, J., Deng, L., Song, S., Zhao, M., Zhang, Y., Wu, S., Wang, G., Zou, Z., Wu, Z., He, W., Chen, F., Deng, N., Wu, S., Wang, Y., Wu, Y., Yang, Z., Ma, C., Li, G., Han, W., Li, H., ... & Shi, L. (2019). Towards artificial general intelligence with hybrid Tianjic chip architecture. *Nature*, 572, 106–111.

Perkins, M. (2023). Academic integrity considerations of AI large language models in the post-pandemic era: ChatGPT and beyond. *Journal of University Teaching & Learning Practice*, *20*(2).

Peterson, L. E. (2009). K-nearest neighbor. *Scholarpedia*, *4*(2), 1883.

Phillips, A. (2021, April 1). *A history and timeline of big data.* TechTarget.

Pitt, C. S., Mulvey, M. S., & Kietzmann, J. (2018). Quantitative insights from online qualitative data: An example from the health care sector. *Psychology & Marketing*, *35*(12), 1010-1017.

Poria, S., Cambria, E., Winterstein, G., & Huang, G. B. (2014). Sentic patterns: Dependency based-rules for concept-level sentiment analysis. *Knowledge-Based Systems*, *69*, 45-63.

Press, G. (2020, April 27). 12 AI milestones: 4. MYCIN, an expert system for infectious disease therapy. *Forbes*.

Qin, X., & Jiang, Z. (2019). The impact of AI on the advertising process: The Chinese experience. *Journal of Advertising*, *48*(4), 338–346.

Qualtrics. (2018). *How AI will reinvent the market research industry*. Qualtrics.

Quijano-Sanchez, L., & Liberatore, F. (2017). The BIG CHASE: A decision support system for client acquisition applied to financial networks. *Decision Support Systems, 98*, 49-58.

Ravi, K., & Ravi, V. (2015). A survey on opinion mining and sentiment analysis: Tasks, approaches and applications. *Knowledge-Based Systems, 89*, 14–46.

Rawat, B., Mehra, N., Bist, A. S., Yusup, M., & Sanjaya, Y. P. A. (2022). Quantum computing and AI: Impacts & possibilities. *ADI Journal on Recent Innovation*, *3*(2), 202-207.

Ray, P. P. (2023). ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet of Things and Cyber-Physical Systems*, *3*, 121–154.

Reisenbichler, M., Reutterer, T., & Schweidel, D. A. (2023). *Applying large language models to sponsored search advertising.* MSI Working Paper Series. MSI.

Reisenbichler, M., Reutterer, T., Schweidel, D. A., & Dan, D. (2022). Frontiers: Supporting content marketing with natural language generation. *Marketing Science, 41*(3), 441–452.

Rialti, R., Marzi, G., Ciappei, C., & Busso, D. (2019). Big data and dynamic capabilities: A bibliometric analysis and systematic literature review. *Management Decision, 57*(8), 2052–2068.

Riikkinen, M., Saarijärvi, H., Sarlin, P., & Lähteenmäki, I. (2018). Using artificial intelligence to create value in insurance. *International Journal of Bank Marketing, 36*(6), 1145-1168.

Rodgers, S. (2021). Themed issue introduction: Promises and perils of artificial intelligence and advertising. *Journal of Advertising, 50*(1), 1-10.

Roy, A., Huh, J., Pfeuffer, A., & Srivastava, J. (2017). Development of trust scores in social media (TSM) algorithm and application to advertising practice and research. *Journal of Advertising, 46*(2), 269–282.

Ruden, J. (2023, October). *Redefining what's possible in paid search performance using AI.* WARC.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature, 323,* 533-536.

Russell, S. (2020). *Human compatible: Artificial intelligence and the problem of control*. Penguin Books.

Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: A modern approach (3rd ed.)*. Pearson Education.

Rust, R. T. (2020). The future of marketing. *International Journal of Research in Marketing, 37*(1), 15-26.

Ryan-Mosley, T., Heikkilä, M., & Yang, Z. (2024, January 5). *What's next for AI regulation in 2024?* MIT Technology Review.

Samuel, A., White, G. R. T., Thomas, R., & Jones, P. (2021). Programmatic advertising: An exegesis of consumer concerns. *Computers in Human Behavior, 116*, 106657.

Sarker, I. H. (2021a). Machine learning: Algorithms, real-world applications and research directions. *SN Computer Science, 2,* 160.

Sarker, I. H. (2021b). Deep learning: a comprehensive overview on techniques, taxonomy, applications and research directions. *SN Computer Science*, *2*, 420.

Sarker, I. H. (2021c). Data science and analytics: an overview from data-driven smart computing, decision-making and applications perspective. *SN Computer Science*, *2*, 377.

Sarker, I. H. (2021d). Deep cybersecurity: a comprehensive overview from neural network and deep learning perspective. *SN Computer Science*, 2, 154.

Sejnowski, T. J., & Rosenberg, C. R. (1988). NETtalk: A parallel network that learns to read aloud. In J. A. Anderson, & E. Rosenfeld (Eds.), *Neurocomputing: foundations of research* (pp. 661-672). MIT Press.

Shahrabi, J., Hadavandi, E., & Asadi, S. (2013). Developing a hybrid intelligent model for forecasting problems: Case study of tourism demand time series. *Knowledge-Based Systems, 43*, 112-122.

Shumanov, M., Cooper, H., & Ewing, M. (2022). Using AI predicted personality to enhance advertising effectiveness. *European Journal of Marketing*, 56(6), 1590-1609.

Sjödin, D. R., Parida, V., Leksell, M., & Petrovic, A. (2018). Smart factory implementation and process innovation. *Research-Technology Management*, *61*(5), 22-31.

Silva, S. C., Duarte, P. & Sundetova, A. (2020), Multichannel versus omnichannel: A price-segmented comparison from the fashion industry. *International Journal of Retail and Distribution Management, 48*(4), 417-430.

Silva, S. C., Martins, C. C., & de Sousa, J. M. (2018). Omnichannel approach: Factors affecting consumer acceptance. *Journal of Marketing Channels, 25*(1-2), 73-84.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*, 484-489.

Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*.

Sohail, S. S., Farhat, F., Himeur, Y., Nadeem, M., Madsen, D. Ø., Singh, Y., Atalla, S., & Mansoor, W. (2023). Decoding ChatGPT: A taxonomy of existing research, current challenges, and possible future directions. *Journal of King Saud University—Computer and Information Sciences*, *35*(8), 101675.

Spatharioti, S. E., Rothschild, D. M., Goldstein, D. G., & Hofman, J. M. (2023). Comparing traditional and LLM-based search for consumer choice: A randomized experiment. *arXiv*:2307.03744.

Spitale, G., Biller-Andorno, N., & Germani, F. (2023). AI model GPT-3 (dis)informs us better than humans. *arXiv:2301.11924*.

Stam, A. (2023, November 1). How AI is expected to reshape the ad industry. *AdAge*.

Stokel-Walker, C., & Van Noorden, R. (2023, February). What ChatGPT and generative AI mean for science. *Nature*, *614*, 214-216.

Stone, M., Aravopoulou, E., Ekinci, Y., Evans, G., Hobbs, M., Labib, A., Laughlin, P., Machtynger, J., & Machtynger, L. (2020). Artificial intelligence (AI) in strategic marketing decision-making: A research agenda. *The Bottom Line*, *33*(2), 183-200.

Sun, S., Cegielski, C. G., Jia, L., & Hall, D. J. (2016). Understanding the factors affecting the organizational adoption of big data. *Journal of Computer Information Systems*, *58*(3), 193–203.

Syam, N., & Sharma, A. (2018). Waiting for a sales renaissance in the fourth industrial revolution: Machine learning and artificial intelligence in sales research and practice. *Industrial Marketing Management*, *69*, 135-146.

Synaptiq. (2020, May 17). From the blog: Artificial intelligence, a brief history. Synaptiq.

Tang, C., & Guo, L. (2015). Digging for gold with a simple tool: Validating text mining in studying electronic word-of-mouth (eWOM) communication. *Marketing Letters*, *26*, 67–80.

Tegmark, M. (2018). *Life 3.0: Being human in the age of artificial intelligence*. Penguin Books.

Terwiesch, C., & Ulrich, K. (2023, September 9). M.B.A. students vs. AI: Who comes up with more innovative ideas? *The Wall Street Journal.*

Tirunillai, S., & Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent Dirichlet allocation. *Journal of Marketing Research*, *51*(4), 463–479.

Tomašev, N., Cornebise, J., Hutter, F., Mohamed, S., Picciariello, A., Connelly, B., Belgrave, D. C. M., Ezer, D., van der Haert, F. C., Mugisha, F., Abila, G., Arai, H., Almiraat, H., Proskurnia, J., Snyder, K., Otake-Matsuura, M., Othman, M., Glasmachers, T., de Wever, W., Teh, Y. W., ..., & Clopath, C. (2020). AI for social good: unlocking the opportunity for positive impact. *Nature Communications*, *11,* 2468.

Tong, S., Luo, X., & Xu, B. (2020). Personalized mobile marketing strategies. *Journal of the Academy of Marketing Science*, *48*, 64-78.

Tong, Y., & Zhang, L. (2023). Discovering the next decade's synthetic biology research trends with ChatGPT. *Synthetic and Systems Biotechnology*, *8*(2): 220-223.

Turing, A. M. (1950). I.—computing machinery and intelligence. *Mind*, 59(236), 433–460.

Vakratsas, D., & Wang, X. (2020). Artificial intelligence in advertising creativity. *Journal of Advertising*, *50*(1), 39–51.

Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, *27*(11), 1134-1142.

Valter, P., Lindgren, P., & Prasad, R. (2018). Advanced business model innovation supported by artificial intelligence and deep learning. *Wireless Personal Communications*, *100*, 97-111.

van Noort, G., Himelboim, I., Martin, J., & Collinger, T. (2020). Introducing a model of automated brand-generated content in an era of computational advertising. *Journal of Advertising*, *49*(4), 411–427.

van Pinxteren, M. M. E., Wetzels, R. W. H., Rüger, J., Pluymaekers, M., & Wetzels, M. (2019). Trust in humanoid robots: Implications for services marketing. *Journal of Services Marketing*, *33*(4), 507–518.

Vlačić, B., Corbo, L., Silva, S. C., & Dabić, M. (2021). The evolving role of artificial intelligence in marketing: A review and research agenda. *Journal of Business Research*, *128*, 187–203.

Vranica, S. (2023, June 19). How AI has the advertising business excited—and worried. *The Wall Street Journal*.

Wang, H., & Hong, W. (2006). Managing customer profitability in a competitive market by continuous data mining. *Industrial Marketing Management*, *35*(6), 715–723.

WARC. (2023, November). *The marketer's toolkit 2024.* WARC.

Wei, C., & Chiu, I. (2002). Turning telecommunications call details to churn prediction: a data mining approach. *Expert Systems with Applications*, *23*(2), 103-112.

Widjaya, I. (2023). *The rise of voice search: How it's changing the digital marketing landscape.* Noobpreneur.

Wirth, N. (2018). Hello marketing, what can artificial intelligence help you with? *International Journal of Market Research*, *60*(5), 435-438.

Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: service robots in the frontline. *Journal of Service Management*, *29*(5), 907-931.

World Health Organization. (2024). WHO.

Wu, C. H., Ho, G. T. S., Lam, C. H. Y., & Ip, W. H. (2015). Franchising decision support system for formulating a center positioning strategy. *Industrial Management and Data Systems*, *115*(5), 853-882.

Xu, F., Uszkoreit, H., Du, Y., Fan, W., Zhao, D., & Zhu, J. (2019). Explainable AI: A brief survey on history, research areas, approaches and challenges. In J. Tang, M. Y. Kan, D. Zhao, S. Li, & H. Zan (Eds.), *Natural Language Processing and Chinese Computing. NLPCC 2019. Lecture Notes in Computer Science, 11839* (pp. 563-574). Springer International Publishing.

Xu, R., & Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Transactions on Neural Networks, 16*(3), 645-678.

Yu, C. (2020). Humanlike robots as employees in the hotel industry: Thematic content analysis of online reviews. *Journal of Hospitality Marketing & Management, 29*(1), 22–38.

Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications, 10*, 3770.

Zhang, H., Rao, H., & Feng, J. (2018). Product innovation based on online review data mining: A case study of Huawei phones. *Electronic Commerce Research, 18*, 3–22.

Zhu, K. (2023, August 3). *The state of state AI laws*: 2023. Epic.org.

Zhu, X. (2005). *Semi-supervised learning literature survey*. University of Wisconsin-Madison Department of Computer Sciences.

Zimand-Sheiner, D., & Earon, A. (2019). Disruptions of account planning in the digital age. *Marketing Intelligence & Planning*, 37(2), 126-139.