Artificial Intelligence and Marketing

ARTIFICIAL INTELLIGENCE AND MARKETING

PIERREYANNDOLBEC

Artificial Intelligence and Marketing Copyright © by pierreyanndolbec. All Rights Reserved.

CONTENTS

Introduction 1

Part I. Main Body

1. Introduction to AI 7
2. Ethics in AI and marketing 58
3. AI and the Transformation of Consumer Experiences 101
4. AI and the Customer Journey 165
5. AI and Competitive and Customer Intelligence 204
6. Generative AI 238
7. AI Agents 267

Appendix 305

Note: This textbook is in beta testing. Use at your own risk.

Artificial Intelligence (AI) is revolutionizing the field of marketing, offering unprecedented opportunities for personalization, efficiency, and customer engagement. This textbook explores the intersection of AI and marketing, providing a comprehensive overview of how AI technologies are reshaping marketing strategies and practices.

This textbook introduces key concepts of AI in marketing. More specifically:

Chapter 1 explains the evolution of Artificial Intelligence from its theoretical beginnings with Alan Turing to its current state as a transformative technology. It details how AI works through data analysis, pattern recognition, and machine learning algorithms. The chapter breaks down core concepts including algorithms, machine learning types, deep learning, natural language processing, and robotics. Finally, it presents two frameworks for categorizing AI: one based on similarity to human behaviors and another based on technological capabilities.

Chapter 2 examines the ethical implications of AI in marketing through five fundamental principles: beneficence, non-maleficence, autonomy, justice, and explicability. It explores how these principles interact and sometimes conflict across different stakeholder levels, while discussing AI's

potential to promote societal good alongside the challenges of balancing effective marketing with environmental responsibility.

Chapter 3 analyzes how AI transforms consumer experiences through four key dimensions: data capture, classification, delegation, and social interaction. The discussion explores how each dimension creates distinct tensions in consumer engagement with AI systems, demonstrating how these experiences can simultaneously empower and constrain, connect and alienate, serve and exploit consumers.

Chapter 4 explores how AI transforms the customer journey through Personalization, Automation, and Prediction (PAP). The chapter analyzes how these capabilities enhance each journey stage: awareness, consideration, purchase, and post-purchase, demonstrating how AI enables more personalized, efficient, and proactive customer experiences.

Chapter 5 examines how AI revolutionizes customer and competitive intelligence through advanced data analysis and pattern recognition. The discussion details how various technologies enhance market understanding, culminating in a comprehensive case study of Starbucks' use of AI for competitive and customer intelligence.

Chapter 6 provides an in-depth look at Large Language Models (LLMs) and generative AI, explaining their capabilities and limitations. It introduces effective prompting techniques and problem formulation strategies, while emphasizing the

importance of understanding AI as a mathematical mirror reflecting patterns rather than truly "thinking."

Chapter 7 defines AI agents as autonomous programs that can perceive, decide, and act toward goals with minimal human intervention. It explains their observe-plan-act cycle, compares two problem-solving approaches (ReAct vs ReWOO), and categorizes five agent types from simple reflex to learning systems. The chapter examines both strengths and limitations of AI agents, discusses common failure modes, and demonstrates their transformative applications in marketing practice and consumer shopping experiences.

PART I
MAIN BODY

1.

INTRODUCTION TO AI

Over the past seventy years, Artificial Intelligence (AI) has evolved from a theoretical concept first explored by Alan Turing into a transformative technology that touches nearly every aspect of our modern world. This chapter offers a comprehensive examination of AI's journey, current state, and future trajectory.

We begin by exploring AI's historical foundations, starting with Turing's groundbreaking question of whether machines could think like humans. From the coining of the term "AI" at Dartmouth College in 1955 through the development of early systems like the Logic Theorist and Perceptron, to today's sophisticated deep learning models, we trace the key innovations that have shaped the field. This evolution reveals not just AI's technical progress, but also how our understanding of machine intelligence has deepened over time.

To build a thorough understanding of AI, we examine it through three distinct lenses: technological, psychological, and cognitive-behavioral. Each perspective illuminates different aspects of artificial intelligence, from its computational foundations to its relationship with human thought processes.

By considering these multiple viewpoints, we can better grasp both AI's current capabilities and its fundamental nature.

The chapter then breaks down how modern AI systems function, exploring core concepts like algorithms, machine learning, deep learning, natural language processing, and robotics. Through practical examples and real-world applications, we demonstrate how these technologies work together to create intelligent systems that can perceive, learn, and act in increasingly sophisticated ways.

Finally, we present two frameworks for categorizing AI: one based on how closely systems mirror human behavior, and another focused on their technological capabilities. These complementary approaches help us understand both the current state of AI technology and its potential future developments, providing a foundation for analyzing AI's ongoing impact across industries and society.

A Brief History of AI

Over the past seventy years, researchers have worked hard to explore if machines can think like humans, a question first posed by Alan Turing. Turing created the famous "Turing test" to see if a computer could mimic human conversation. This test laid the groundwork for what we now call artificial intelligence (AI). The term AI was first used in 1955 during a summer research workshop at Dartmouth College. AI is

defined as making a machine act in ways that would be considered intelligent if a human were doing the same. In the late 1950s, AI research focused on teaching computers to solve math and logic problems, leading to developments like the Logic Theorist, the Perceptron, and object recognition systems.

Since the 1980s, AI has grown into a booming industry, especially as researchers shifted their focus to machine learning. Machine learning is a part of AI where a computer program learns from experience to perform tasks better over time. It has revolutionized many fields, including biology, education, engineering, finance, healthcare, and marketing. Machine learning algorithms use large amounts of data to find patterns and make predictions, recommendations, and decisions. For example, companies like Amazon and Netflix use machine learning to power their recommendation systems.

The latest advancements in AI are due to deep learning. Deep learning involves using neural networks with many hidden layers to process data in complex ways. This method allows algorithms to refine themselves with new data, similar to how the human brain works. Deep learning has significantly improved AI's ability to make accurate predictions and is used in areas like image recognition, social media analysis, and autonomous driving.

The three key milestones in the evolution of AI are illustrated in Figure 1.1.

Figure 1.1. The evolution of AI.

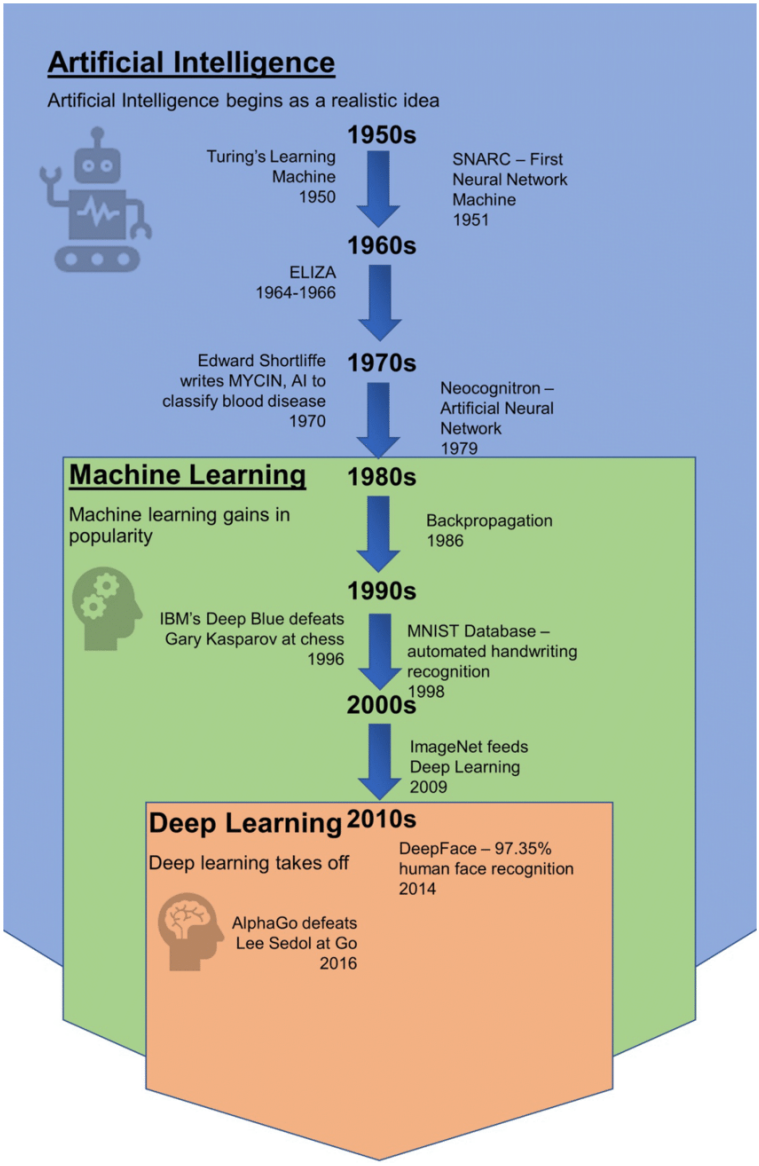

Three Historical Perspectives

on Understanding AI

_*]:min-w-0 !gap-3.5">

Following the brief historical overview, we can trace how our understanding of artificial intelligence has evolved through three major perspectives. Each represents a different era and approach to conceptualizing AI, and together they show how the field has matured from a narrow technical endeavor to a rich interdisciplinary pursuit.

1. AI as a Computational Technology (1950s-1960s)

The earliest AI pioneers approached artificial intelligence primarily as a computational and engineering challenge. This perspective dominated the field's founding years and established AI as a technical discipline.

John McCarthy, who coined the term "artificial intelligence" in 1956, exemplified this view when he defined AI as "the science and engineering of making intelligent machines, especially intelligent computer programs." This definition deliberately separated AI from psychology and biology—it was about building intelligent systems, not necessarily understanding or replicating human cognition.

From this technological perspective, intelligence was reconceptualized as information processing. Key characteristics of this view included:

- **Problem-solving focus**: Early AI concentrated on well-defined problems with clear rules—chess, theorem proving, and mathematical puzzles. The Logic Theorist (1956), often considered the first AI program, proved mathematical theorems by manipulating symbols according to logical rules.
- **Symbol manipulation**: Intelligence was seen as the ability to manipulate symbolic representations according to formal rules. Allen Newell and Herbert Simon's "Physical Symbol System Hypothesis" (1976) argued that symbol manipulation was both necessary and sufficient for intelligence.
- **Goal-oriented behavior**: Richard Sutton and others defined intelligence functionally as "the ability to achieve goals in the world" through computational processes. What mattered wasn't how thinking happened, but whether machines could produce intelligent outputs.
- **Flexibility and adaptation**: Even early researchers recognized that intelligent systems needed to handle uncertainty and incomplete information. Unlike traditional algorithms that required complete data, AI systems needed to make reasonable decisions with partial information.

This technological perspective gave us foundational concepts still used today: search algorithms, knowledge representation,

planning systems, and logical reasoning. However, by the 1970s, researchers realized that many aspects of intelligence—like vision, language understanding, and common-sense reasoning—were far harder than initially anticipated.

2. AI as Cognitive Simulation (1970s-1990s)

As AI faced its first "winter" in the 1970s, researchers began drawing more heavily from psychology and cognitive science. This shift represented a fundamental change: instead of just building intelligent machines, AI researchers started asking how human intelligence actually works.

This cognitive perspective viewed AI as the computational modeling of human mental processes. Key developments included:

- **Cognitive architectures**: Systems like ACT-R (Adaptive Control of Thought-Rational) and SOAR attempted to model the complete structure of human cognition, including memory systems, attention, and learning mechanisms.
- **Mental functions as modules**: AI began breaking down intelligence into specific cognitive capabilities:
    - **Perception**: Pattern recognition, computer vision (like David Marr's computational theory of vision)

    - **Memory**: Short-term vs. long-term storage, semantic networks, frames
    - **Reasoning**: Mental models, analogical reasoning, case-based reasoning
    - **Problem-solving**: Means-ends analysis, heuristic search
    - **Learning**: Concept formation, explanation-based learning
- **Knowledge representation**: Researchers developed sophisticated ways to encode human knowledge, including semantic networks, frames, scripts, and ontologies. The CYC project (started 1984) attempted to encode all human common-sense knowledge.
- **Expert systems**: These systems captured human expertise in narrow domains. MYCIN (1976) diagnosed blood infections, while DENDRAL (1965) identified molecular structures. They represented AI's attempt to model how human experts think and reason.

This era established that intelligence wasn't just about processing power or clever algorithms—it required vast amounts of knowledge and sophisticated ways of organizing and using that knowledge. The cognitive perspective also highlighted the importance of learning, leading to the machine learning revolution of the 1990s.

3. AI as Intelligent Behavior (1990s-Present)

The modern perspective, crystallized in Russell and Norvig's influential textbook (1995), shifts focus from internal processes to external behavior and performance. This behavioral perspective asks not "Does it think like a human?" but "Does it behave intelligently?"

This perspective recognizes four different approaches to AI, based on two dimensions:

The Behavioral Framework

Dimension 1: Thinking vs. Acting

- Thinking: Focus on reasoning and internal cognitive processes
- Acting: Focus on behavior and external performance

Dimension 2: Humanly vs. Rationally

- Humanly: Mimicking human performance (including errors and limitations)
- Rationally: Optimal performance according to logic and probability

This creates four approaches:

1. **Acting Humanly**: The Turing Test approach

    - Goal: Fool humans into thinking the machine is human
    - Success: Natural language processing, knowledge representation, automated reasoning, machine learning
    - Example: Modern chatbots like ChatGPT that can maintain convincing conversations

2. **Thinking Humanly**: Cognitive modeling approach

    - Goal: Program computers to think like humans
    - Method: Use psychological experiments to develop precise theories of mind
    - Example: Neural networks inspired by brain structure

3. **Thinking Rationally**: The "laws of thought" approach

    - Goal: Build systems that use formal logic to derive conclusions
    - Challenge: Not all intelligent behavior is logical; humans often act on intuition
    - Example: Automated theorem provers, formal verification systems

4. **Acting Rationally**: The rational agent approach

    - Goal: Build agents that achieve the best expected outcome
    - Advantage: More general than "laws of thought" – can include reflexes, intuition
    - Example: Recommendation systems, autonomous vehicles, game-playing AI

Why This Shift Matters

The behavioral perspective has several advantages:

- **Practical focus**: It emphasizes what AI can do rather than how closely it mimics human cognition
- **Measurable success**: Performance can be objectively evaluated
- **Flexibility**: Different approaches suit different problems
- **Includes non-human intelligence**: Recognizes that optimal intelligence might differ from human intelligence

This perspective dominates modern AI. Today's systems—from AlphaGo to GPT models—are evaluated primarily on their performance, not their similarity to human cognition. Machine learning systems find patterns humans

might never discover. Computer vision systems can exceed human accuracy. Translation systems work through statistical patterns rather than linguistic understanding.

The Evolution and Integration of Perspectives

These three perspectives aren't mutually exclusive—they represent an evolution and broadening of how we understand AI:

1. **1950s-1960s**: AI as computational technology focused on symbol manipulation and problem-solving
2. **1970s-1990s**: AI as cognitive simulation tried to replicate human mental processes
3. **1990s-Present**: AI as intelligent behavior emphasizes performance over process

Modern AI integrates all three:

- We use **computational techniques** (perspective 1)
- Inspired by **cognitive science** (perspective 2)
- To create systems that **behave intelligently** (perspective 3)

For example, a modern language model like GPT:

- Uses computational technology (transformer architectures, massive parallel processing)
- Incorporates cognitive insights (attention mechanisms, context understanding)
- Is evaluated on behavioral performance (conversation quality, task completion)

This evolution shows how AI has matured from a narrow technical field to a rich interdisciplinary endeavor that draws from computer science, psychology, neuroscience, philosophy, and mathematics. Understanding these perspectives helps us appreciate both where AI came from and where it might be heading.

Types of AI: Two categorizations

Having explored how our conceptual understanding of AI has evolved through three historical perspectives—from computational technology to cognitive simulation to intelligent behavior—we now turn to how we classify and categorize actual AI systems today. While the historical perspectives helped us understand what AI *is* and how we think about it, we need practical frameworks for categorizing the AI systems we build and encounter.

Modern AI researchers and practitioners use two primary classification systems, each serving different purposes:

1. **Behavioral-Cognitive Classification**: Categorizing AI by how closely it mirrors human cognitive capabilities
2. **Capability Scope Classification**: Categorizing AI by the breadth and power of its abilities

These aren't competing systems but complementary frameworks. Just as a biologist might classify animals by their evolutionary relationships (phylogeny) or by their ecological roles (functional groups), we classify AI systems in multiple ways to understand different aspects of their nature and potential.

Why We Need Classification Systems

The historical perspectives we examined show how our understanding of AI has evolved, but they don't give us practical tools for categorizing the AI systems we encounter daily. Consider these questions:

- How do we distinguish between Siri and GPT-4?
- What separates a chess-playing program from a self-driving car's AI?
- How do we prepare for AI systems that don't yet exist?

Classification systems help us:

- **Understand current capabilities**: What can today's AI actually do?
- **Identify limitations**: What can't current systems accomplish?
- **Plan for the future**: What developments should we anticipate?
- **Make strategic decisions**: Which type of AI suits our needs?

Let's explore these two classification systems in detail.

Types of AI according to the similarity with human behaviors

Artificial Intelligence (AI) has evolved significantly, with various types of AI systems developed to mimic human intelligence and behavior to different degrees. This categorization of AI based on similarity to human behaviors provides a framework for understanding the capabilities and limitations of different AI systems.

From the simplest reactive machines to the theoretical concept of self-aware AI, each type represents a step forward in AI's ability to process information, learn, and interact with its environment. This progression not only showcases the

advancements in AI technology but also highlights the ongoing challenges and future possibilities in the field.

In this overview, we'll explore four main types of AI based on their similarity to human behaviors:

1. Reactive Machines
2. Limited Memory
3. Theory of Mind
4. Self-Aware AI

Each type offers unique capabilities and applications, particularly in the field of marketing, where AI is increasingly being used to enhance customer experiences, optimize strategies, and drive business outcomes. Let's delve into each type to understand their characteristics, applications, and implications for the future of AI and marketing.

Reactive Machines

Reactive machines are the simplest form of AI, designed with limited capabilities and functions. They have no internal memory and cannot learn from past experiences. These machines react only to immediate stimuli in their environment, making them task-specific. An algorithm developed for one purpose in a reactive machine cannot be adapted for any other task. Reactive machines lack memory-based functionality and do not have a representation of their

environment. This means they cannot improve over time or remember goals needed to achieve an objective. Their knowledge is limited to a "stimulus-response" (S-R) model, which operates on the principle of "if X is sensed, then do Y." Due to this, reactive machines are often referred to as "S-R agents."

A basic email auto-responder AI reacts to certain triggers in customer emails. If a customer sends an email with the subject "Order Status" (stimulus), the AI responds with a preset message, such as "Thank you for your inquiry. Your order is being processed" (response). It does not remember the customer's previous interactions or modify its responses—it operates purely based on the current message, using a simple stimulus-response framework.

Reactive machines act based on what they see without learning from past events or anticipating future ones. If the input remains the same, their output will always be consistent. They are effective at performing specific tasks but cannot be easily adapted to other tasks or situations. They lack an understanding of the broader world and cannot handle tasks beyond their assigned functions. Despite their limitations, reactive machines are still quite useful. One of their main advantages is predictability, which ensures a level of certainty. This predictability is valuable when safety and reliability are priorities in AI-powered machines. For example, IBM's Deep Blue, a chess-playing computer developed in the 1980s, is a well-known reactive machine. Deep Blue defeated world

champion Garry Kasparov by identifying chess pieces, tracking their movements, predicting future moves, and selecting the optimal one. However, Deep Blue couldn't process past data; it only considered the current positions of the chess pieces and chose from the possible next moves. This example shows how reactive machines can achieve significant improvements in specific functions even without learning or adapting.

Limitations and Applications in Marketing

Reactive machines can only perform the tasks they are programmed for and cannot handle complex, multi-step processes or adapt to new tasks. This makes them less flexible than more advanced AI systems. In marketing, reactive machines can be used for specific, repetitive tasks such as customer service chatbots that provide predefined answers to frequently asked questions.

The predictability of reactive machines is an advantage in scenarios where consistent performance is crucial. For example, in automated customer support, predictable responses ensure that customers receive accurate information. They are used in scenarios requiring high reliability, such as automated email responses, where consistency and correctness are more important than adaptability.

Strategic Implications

- **Cheaper Implementation:** Reactive machines are

typically cheaper to implement and maintain due to their simplicity, making them ideal for straightforward tasks that do not require complex processing. This efficiency and cost-effectiveness can be particularly useful for small businesses or for specific functions within larger organizations where advanced AI might be overkill.

- **Increased Consistency and Reliability:** In terms of customer interaction and experience, reactive machines can provide consistent, reliable interactions but may fall short in delivering personalized or nuanced experiences. They are best suited for initial customer contact points or handling routine inquiries, while more advanced AI can be reserved for complex customer service needs.

Examples in Marketing

- In marketing, reactive machines find their niche in applications where specific and repetitive tasks are necessary. One example is automated customer support. Simple chatbots provide answers to common questions, like order status inquiries or store hours, without the ability to learn or adapt to new questions.
- Another example is programmatic advertising. Algorithms place ads based on predefined rules and criteria, such as displaying a specific ad when a certain keyword is detected on a webpage. These systems react

to immediate data inputs (keywords) without learning from past ad placements.

Limited Memory

Limited memory is a type of machine learning used widely today. It refers to machines that can analyze data, learn from historical inputs or actions, and build knowledge from those experiences. Unlike reactive machines, limited memory systems not only react to the current environment but also improve their performance over time based on past data and experiences. This kind of system needs training data and observations to learn and make better decisions or predictions. Types of learning under limited memory include reinforcement learning and supervised/unsupervised learning.

A common example of limited memory is image recognition. Systems are trained with thousands or even millions of labeled pictures to recognize and identify objects accurately. The more images the system is trained on, the better it becomes at recognizing new images. Another example is self-driving cars, which use recent data like speed, speed limits, and the distance from other vehicles to make real-time decisions and avoid accidents. Chatbots, which improve customer service by providing quick responses and guiding users to the correct services, also rely on limited memory.

Despite its advantages over reactive machines, limited memory has some limitations. It requires vast amounts of data

to learn even simple tasks. Moreover, it can only retain information for a short period. Unlike humans, these systems cannot store long-term memories, meaning they must be retrained from scratch if the environment changes. For instance, self-driving cars retain memory only for the duration of a journey and cannot build on past experiences to improve driving skills over time. Humans, by contrast, can transfer knowledge and skills from one area to another more easily.

Applications in Marketing

Among the various types of AI, limited memory is the most commonly used in marketing. It employs machine learning techniques like deep neural networks and ensemble trees to predict future demand and understand customer behavior. These techniques allow marketers to gather information efficiently and respond quickly to customer needs. Big data analytics, powered by limited memory's data mining techniques, provides valuable consumer insights.

For example, chatbots and voice bots are transforming customer service by providing quick, accurate responses and guiding users to the correct services based on past interactions. Virtual assistants in self-driving cars use recent data to provide navigation assistance and improve the driving experience, showcasing how limited memory AI can enhance user interaction.

Strategic Implications

- **Efficiency and Personalization:** Limited memory systems allow for efficient data processing and personalized marketing efforts. By analyzing past consumer behavior, marketers can tailor their strategies to individual preferences, increasing engagement and conversion rates.
- **Predictive Analytics:** These systems are crucial for predictive analytics, helping businesses forecast trends and consumer demands. This enables companies to adjust their strategies proactively, ensuring they stay ahead of the competition.

Examples in Marketing

- **Recommendation Systems:** Online retailers like Amazon use limited memory AI to recommend products to customers based on their browsing history and previous purchases. This not only enhances the shopping experience but also boosts sales by suggesting relevant items.
- **Dynamic Pricing: A**irlines and hotels use machine learning algorithms to adjust prices in real-time based on demand, competition, and other factors. Limited memory AI helps analyze historical pricing data and current market

conditions to set optimal prices.

Comparison with Reactive Machines

Limited memory AI systems are a step above reactive machines as they can learn from historical data and improve over time. While reactive machines are confined to predefined tasks and immediate stimuli, limited memory systems adapt and evolve, making them more versatile and useful in dynamic environments like marketing. Take, for example, chatbots from the two types of AI systems:

- **Reactive Machine Chatbot:** A reactive machine chatbot is a type of AI that responds to specific prompts or inputs but does not have the ability to store past interactions or learn from them. It operates purely based on predefined rules or responses, reacting to user inputs without any memory of previous interactions. For example, a simple FAQ bot that answers questions based on a fixed set of rules does not recall previous conversations. Reactive machines are limited to predefined tasks and immediate stimuli, making them less adaptable.
- **Limited Memory Chatbot:** In contrast, a limited memory chatbot can remember past interactions within a session or over a short period, using this memory to provide more contextually relevant responses. These chatbots store some information temporarily to improve

user experience, such as remembering the user's name or previous questions during a session. Although they do not have long-term memory capabilities, limited memory chatbots can adapt and evolve based on recent interaction history, making them more versatile and useful in dynamic environments like marketing. For example, a customer service bot that remembers details about your issue during a single session but forgets once the session ends demonstrates this capability.

Theory of Mind

Theory of Mind (ToM) in AI is a stage where systems can interact with the world similarly to humans, considering needs, emotions, beliefs, and thought processes. Unlike other types of AI, ToM accounts for human emotions. Traditional machine learning can recognize human actions but struggles when people change their minds due to their mental states and thoughts. This gap highlighted the need for AI that understands how emotions influence human actions and thoughts, leading to the development of ToM.

ToM involves discussions about the mind and mental states. In psychology, ToM is the ability to attribute mental states to oneself and others. Premack and Woodruff introduced this concept in 1978 to explore how young children understand their own and others' mental states. The theory revealed gaps in perceptions, knowledge, and intentions between ourselves

and others, which affect social interactions. Human social skills develop by attributing our desires, beliefs, and goals to others. Based on this theory, AI researchers aim to create machines capable of socializing in ways that weren't possible before.

The idea of socially intelligent robots is widely discussed. Most researchers agree that ToM should enable machines to gather meaningful information from their social environment in real-time and respond with human-like behavior. These robots should observe and learn from humans, similar to how children do. They should internalize information, ask questions, form opinions, and make decisions autonomously. Robots should also develop social cues and express emotions and desires through interactions without relying solely on pre-programmed vocabulary. Additionally, they should recognize these states in others and act accordingly.

Strategic Implications

- **Enhanced Consumer Insights:** ToM AI provides deeper insights into consumer behavior by analyzing not just actions but the underlying emotional and psychological drivers.

- **Human-Like Interactions:** The ability of ToM AI to mimic human social interactions makes it invaluable in roles requiring empathy and nuanced

communication.

Applications in Marketing

In marketing, the development of ToM AI holds significant promise. By understanding and predicting human emotions, needs, and thought processes, these advanced AI systems can create more personalized and emotionally resonant marketing strategies.

- **Personalized Customer Interactions:** ToM AI can enhance customer service by tailoring interactions based on the customer's emotional state and preferences. This leads to more meaningful and satisfying customer experiences.

Example: Virtual shopping assistants that can detect frustration or confusion in a customer's voice and adjust their responses to provide better assistance.

- **Emotion-Driven Marketing Campaigns:** Understanding consumer emotions allows marketers to craft campaigns that resonate on a deeper emotional level, increasing engagement and effectiveness.

Example: Ad campaigns that adapt in real-time based on the

viewer's emotional reactions, leading to higher conversion rates.

Comparison with Other AI Types

- **ToM AI vs. Reactive Machines:** While reactive machines operate on predefined rules and stimuli-response mechanisms, ToM AI can understand and react to emotional and cognitive states, making it more sophisticated and versatile.
- **ToM AI vs. Limited Memory AI:** Limited memory AI systems learn from past data but do not account for real-time emotional and psychological states. ToM AI bridges this gap by incorporating an understanding of human mental states into its decision-making processes.

Future Potential

Social robots, due to their human-like advancements, are in demand in areas like personal assistance, senior care, education, and public services. However, ToM is complex and still developing, with a long way to go before it achieves substantial and meaningful realization. Sophia, a social humanoid robot, exemplifies efforts in this direction. Sophia can recognize individuals, follow faces, and socialize using advanced vision, machine learning, artificial intelligence, animatronics, and natural language processing. However, some critics view Sophia as more of a puppet than a truly

intelligent being. Despite the controversy, Sophia represents a step towards superintelligence.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=136#oembed-1

Understanding these aspects helps marketers leverage the strengths of ToM AI while being aware of its limitations, allowing for strategic deployment in marketing operations. This includes using AI to enhance customer experiences, predict market trends, and personalize marketing efforts, ultimately driving better business outcomes.

Self-Aware AI

The most advanced level of artificial intelligence (AI) is self-aware AI. This type of AI would evolve to the point where it can develop self-awareness, similar to humans. Self-awareness refers to being consciously aware of one's internal states and interactions with others. In AI, a machine is considered self-aware if it can think, feel, and understand its own emotions. Such machines would not only interact with the world and

learn from it but also form representations of themselves. While this is still a theoretical concept, many AI scientists believe that creating self-aware AI is important. Self-aware systems could better manage their goals and complex environments.

AI scientists have been striving to build highly autonomous systems that can manage themselves. These advanced machines would need to be able to run independently, adjust to different situations, and handle workloads efficiently. They should understand user demands and manage tasks without needing user intervention. Furthermore, autonomous systems would have "self-healing" abilities, allowing them to detect, diagnose, and repair problems automatically. They could protect themselves against attacks or failures, learn from these issues, and prevent future problems.

Achieving self-aware AI requires a diverse and comprehensive understanding of intelligence. Success in creating such a system involves integrating cognition (planning, understanding, and learning) with metacognition (control and monitoring of cognition) and intelligent behaviors. AI needs to be adaptable and capable of switching methods if one approach fails. Knowledge from various fields, including engineering, neuroscience, and psychology, is crucial for developing self-aware AI.

Future Potential

Self-aware AI, although still theoretical, represents the

pinnacle of AI development with profound implications for various industries, including marketing. While true self-aware machines do not yet exist, ongoing research continues to push the boundaries of what AI can achieve. For instance, efforts to develop robots like Sophia aim to incorporate elements of self-awareness, empathy, and social interaction, albeit on a rudimentary level.

Understanding these aspects helps marketers leverage the strengths of advanced AI while being aware of its limitations, allowing for strategic deployment in marketing operations. This includes using AI to enhance customer experiences, predict market trends, and personalize marketing efforts, ultimately driving better business outcomes.

Types of AI Based on Technology Levels

While the behavioral-cognitive classification focuses on how AI processes information, this system—popularized by philosopher Nick Bostrom and others—categorizes AI by the breadth of its capabilities. It answers: "What range of tasks can this AI perform?"

In this overview, we'll examine three main types of AI based on their technological sophistication:

1. Artificial Narrow Intelligence (ANI)
2. Artificial General Intelligence (AGI)

3. Artificial Super Intelligence (ASI)

Each of these categories represents a significant leap in AI capabilities, with implications that extend across various industries, including marketing. As we explore each type, we'll consider their characteristics, applications (both current and potential), and the impact they may have on business strategies and consumer interactions.

This alternative categorization complements our understanding of AI's relationship to human behavior, offering a more complete picture of the AI landscape and its future possibilities. Let's delve into each type to grasp the full spectrum of AI technology and its evolving role in shaping our world.

Artificial Narrow Intelligence (ANI)

Also known as "weak AI," Artificial Narrow Intelligence (ANI) refers to systems designed to handle a specific task. For example, voice assistants like Siri and Alexa can recognize speech and respond with relevant information, but they can't perform tasks outside their programmed functions. ANI refers to AI designed to perform a specific task autonomously using human-like intelligence. According to Kurzweil (2005), narrow AI consists of systems that carry out specific intelligent behaviors in particular contexts. The term "narrow" indicates its limited capabilities. If the input changes, the system needs

reprogramming to maintain its performance (Goertzel, 2014). ANI can only handle tasks it has been pre-programmed for. Unlike human intelligence, which can adapt to environmental changes, ANI cannot. It mimics human behavior based on given goals, contexts, and data but cannot transfer knowledge between different tasks as humans can (Taylor et al., 2008). Reactive machines and systems with limited memory fall into this category. Most current AI systems are ANI.

Artificial General Intelligence (AGI)

Known as "strong AI," AGI refers to machines that possess the ability to understand, learn, and apply knowledge in ways that are comparable to human intelligence. Unlike ANI, AGI systems could potentially perform any intellectual task a human can do. Imagine a robot that can not only converse with you but also understand emotions, solve complex problems, and adapt to new situations just like a human being. It is a concept with various definitions and no clear consensus. It can refer to a system's capabilities, the pursuit of creating such systems, or the study of their nature. Despite the lack of a precise definition, AGI is generally understood as possessing human-like intelligence.

Most researchers agree on several key characteristics of AGI:

- **Versatility:** AGI can achieve diverse goals and handle various tasks in complex and changing environments.

For instance, an AGI system could excel at both driving a car and diagnosing medical conditions.

- **Knowledge Generalization:** AGI should be able to learn from previous experiences and apply that knowledge to different contexts. For example, an AGI that learns to play one video game should be able to transfer strategies and skills to another game.
- **Autonomous Problem Solving:** AGI is expected to solve problems independently, beyond human inputs and expectations. This means an AGI could potentially discover new scientific theories or invent novel technologies without direct human guidance.

These attributes suggest that AGI would be capable of learning, understanding, and performing intellectual tasks much like a human. The goal is to build a system with a broad and adaptable scope of intelligence.

The idea of AGI aligns with the early ambitions of AI's founding figures like John McCarthy, Marvin Minsky, Allen Newell, and Herbert A. Simon, who envisioned creating systems comparable to the human brain. However, as progress stalled, researchers shifted focus to more achievable and specific problems, moving away from the grand vision of AGI.

Currently, achieving AGI is challenging due to limitations in techniques and resources. Despite skepticism, many scientists remain hopeful about the possibility of AGI. Researchers like Wang and Goertzel have indicated that with

continued advancements, an AGI system could become a reality in the future.

Future Potential

Artificial General Intelligence (AGI), although still theoretical, would have profound implications for various industries, including marketing. While true AGI has not yet been realized, ongoing research continues to push the boundaries of what AI can achieve. For instance, advanced AI systems like DeepMind's AlphaGo and AlphaZero demonstrate early signs of generalized learning and problem-solving.

Understanding these aspects helps marketers prepare for the transformative potential of AGI while being aware of its limitations, allowing for strategic planning and adaptation to future advancements in AI technology. This includes leveraging AGI to enhance consumer insights, personalize marketing efforts, and drive innovation in marketing strategies, ultimately leading to more effective and efficient marketing operations.

Artificial Super Intelligence (ASI)

ASI is another hypothetical form of AI that surpasses human intelligence in all aspects, including creativity, problem-solving, and emotional intelligence. This means that ASI would not only match the abilities of Artificial General

Intelligence (AGI) but also significantly outperform human beings in tasks due to its superior memory, efficient data processing, and advanced analytical skills. Unlike the human brain, which is limited by biological constraints, ASI would operate on a computer substrate, offering benefits like modularity, transparency in operations, and superior computation abilities.

ASI, also known as transhuman AI, represents a stage where humans can transcend their natural limitations through AI technologies. Many scientists believe ASI could solve various human problems, including aging and incurable diseases. Some suggested that ASI would be more efficient, unbiased, resourceful, and continuously available, potentially even making human immortality achievable by transferring human consciousness into computer systems.

However, there are significant social and ethical concerns regarding ASI. Some emphasized that superintelligence should be developed for the benefit of all humanity and aligned with shared ethical principles. To mitigate fears of autonomous weapons and other potential dangers, researchers propose defining and embedding minimal standards of ethics into superintelligent AI systems. This includes setting goals that ensure the AI acts "friendly" and does not harm humans. For example, an AI could be given a set of ethical guidelines based on historical figures known for their morality, instructing it to avoid actions these individuals would disapprove of.

In essence, while ASI holds the promise of solving many

human problems, it also presents potential risks. It is up to humanity to ensure that these risks are managed through ethical guidelines and common sense, maintaining control over the development and application of superintelligent AI.

How Does AI Work?

At its core, AI revolves around the use of data to learn and improve. By analyzing vast amounts of data, AI systems can identify patterns and relationships that might be difficult for humans to discern.

The learning process of AI typically involves algorithms, which are sets of rules or instructions that guide the AI in analyzing data and making decisions. In machine learning, a widely used branch of AI, algorithms are trained on datasets that can be either labeled (where the data is tagged with the correct answer) or unlabeled (where the AI finds patterns without explicit guidance). This training enables the AI to make predictions or categorize new information.

Deep learning, a more advanced subset of machine learning, uses artificial neural networks composed of multiple layers. These neural networks function in a way that resembles the human brain, allowing the AI to process information at different levels of abstraction. This layered approach helps AI systems learn complex tasks, such as recognizing objects in

images, understanding spoken language, and translating text between languages.

As AI systems are exposed to more data and experience, they continuously learn and adapt, becoming better at the tasks they are designed to perform. For instance, AI is now used in applications ranging from personalized marketing recommendations on platforms like Amazon and Netflix to advanced language translation services such as Google Translate.

Basic Concepts of AI

Algorithms

In computer science, algorithms are a series of instructions or steps that computer programs follow to complete a specific task, such as turning data into useful information. An algorithm is any well-defined procedure that takes some values as input and produces some values as output. For example, an algorithm could be used to sort a list of numbers or find the shortest path between two points on a map.

A great real-life example of an algorithm is **Google Maps**. When you enter a starting point and destination, Google Maps uses algorithms to calculate the best route. The underlying algorithm, known as **Dijkstra's algorithm** or its variations, finds the shortest or fastest path between two points on a map.

Machine Learning (ML)

Machine learning is a vast and rapidly evolving field where computers learn from data to make decisions with minimal human intervention. ML systems can perform a wide range of complex tasks, such as building search engines, filtering content on social media, making recommendations on websites, predicting traffic patterns, trading stocks, and diagnosing medical conditions. You can learn more about the basics of how machine learning operates here.

There are six main types of ML tasks:

- **Supervised Learning:** In supervised learning, a model is trained using a labeled dataset that includes both input variables (X) and output variables (Y). The goal is to learn a function that can predict the output based on new inputs accurately.

Example: Predicting house prices based on features like size, location, and number of bedrooms.

- **Unsupervised Learning:** In unsupervised learning, a model is training using unlabeled datasets and only input variables, and the goal is to find hidden patterns or extract useful information. Common tasks include clustering, where data points are grouped based on similarity, and dimensionality reduction, where high-

dimensional data is simplified while retaining essential information.

Clustering example: Spotify uses clustering to group listeners based on their music preferences and listening habits. These clusters help create personalized playlists like "Discover Weekly" by grouping users with similar tastes.

Dimensionality reduction example: Facial Recognition: Social media platforms like Facebook use dimensionality reduction to simplify facial image data by reducing thousands of pixel values to key facial features. This allows for faster and more accurate matching of faces in photos.

- **Semi-Supervised Learning:** Semi-supervised learning combines a small amount of labeled data with a large amount of unlabeled data to improve model accuracy. This is particularly useful when labeled data is expensive or time-consuming to obtain.

Example: Google Translate uses semi-supervised learning to improve the quality of its translations. Initially, a small dataset of human-translated sentence pairs (labeled data) is used to train the model. This is combined with a large amount of monolingual text in different languages (unlabeled data) to further refine the language model. By leveraging the patterns and structures found in the unlabeled text, the system can

learn to improve translations for languages with limited labeled datasets.

- **Transfer learning:** Transfer learning involves taking an existing model trained on one task and adapting it for a different but related task. This is useful when there is limited data for the new task but ample data for a similar task.

Example: Google Health's AI system for breast cancer detection leveraged transfer learning by starting with a pre-trained model built on non-medical image data and fine-tuning it with a smaller dataset of annotated mammograms. This approach allowed the system to identify subtle patterns in medical images with high accuracy, despite the relatively limited availability of labeled medical data.

- **Active Learning:** In active learning, the model actively selects the most informative data for training to optimize learning with minimal data acquisition. This approach is ideal when gathering new data is costly or time-consuming. The goal is to maximize predictive accuracy with the minimum amount of data.

Example: Pfizer might train a machine learning model on a small set of known chemical compounds labeled as effective or ineffective for a specific disease. The active learning algorithm

then identifies the most uncertain or informative compounds from a large pool of unlabeled candidates. These selected compounds are tested in the lab, and the results are fed back into the model to improve its predictions iteratively.

- **Reinforcement Learning:** In reinforcement learning, an agent interacts with its environment, taking actions and receiving feedback to optimize its performance. The agent learns to achieve a goal by maximizing cumulative rewards. This approach is used in applications like autonomous vehicles, where the system must navigate and make decisions based on real-time data, and in game playing, where the agent learns strategies to win games.

Example: Waymo, a self-driving car company, uses reinforcement learning to train its autonomous vehicles to navigate roads safely and efficiently. The vehicle (agent) interacts with its environment (real-world or simulated driving conditions) by taking actions such as steering, accelerating, or braking.

The system receives *feedback* in the form of rewards or penalties:

- *Rewards* for successful actions (e.g., maintaining a safe distance from other vehicles, following traffic rules).
- *Penalties* for undesirable outcomes (e.g., near misses, collisions, or breaking traffic laws).

Over time, the reinforcement learning model learns to optimize its behavior by maximizing cumulative rewards. This enables the vehicle to make complex real-time decisions, such as merging into traffic, navigating intersections, and handling unexpected obstacles, all while ensuring passenger safety and comfort.

Deep Learning

Deep learning is a subset of ML that uses artificial neural networks with many layers (hence "deep") to process large amounts of data and learn complex patterns. These networks are designed to mimic how the human brain works, with layers of interconnected nodes (neurons) that process information. Deep learning can perform tasks like image recognition, speech recognition, and natural language processing. You can learn more about how deep learning works here.

- **Image Recognition:** Deep learning systems can identify objects, people, places, and other elements in images. For example, Facebook's automatic friend tagging uses deep learning to suggest tags for people in photos by recognizing their faces.
- **Speech Recognition:** Systems like Google Voice, Siri, and Alexa use deep learning to understand and respond to spoken commands. They convert spoken language into text and process the information to generate

responses, enabling tasks like voice search, setting reminders, and controlling smart home devices.

Natural Language Processing (NLP)

NLP combines AI and linguistics to enable computers to understand and generate human language. NLP systems can perform tasks such as automatic language translation, information retrieval, text summarization, question answering, and creating chatbots. For example, Google Translate uses NLP to translate text between languages, and chatbots use NLP to interact with users and answer their questions. A major challenge in NLP is the lack of resources and tools for many low-resource languages spoken by millions of people, while most resources are focused on high-resource languages like English and Mandarin. You can learn more about how large language models work here.

For example, ChatGPT uses NLP to analyze and generate human-like responses to user queries, making it effective for tasks like customer support and content creation. It processes input text, identifies intent, and crafts relevant answers, showcasing the power of conversational AI.

Robotics

AI has been closely connected with robotics since its early days. The goal is not only to build robots that can act like humans

but also to think like them. Robotics involves the intelligent connection of perception to action. Modern robots, integrated with AI, can recognize, reason, and act based on sensor inputs. These intelligent robots can perform tasks such as navigating spaces, avoiding obstacles, and manipulating objects. For example, robots in manufacturing can assemble products with precision, and autonomous drones can navigate and deliver packages. Future developments aim to create robots that can operate in a wide range of challenging environments and assist in various tasks, potentially even serving as companions or aides for humans.

With the integration of AI, robots are no longer just machines that perform manual tasks. They are now equipped with sophisticated software that allows them to make decisions and take actions based on their surroundings. For instance, self-driving cars use a combination of sensors, cameras, and AI algorithms to navigate roads and avoid obstacles. AI-powered robots can also learn from their experiences, improving their performance over time.

Key Terms

Category	Term	Definition
Core Concepts	**Artificial Intelligence (AI)**	The field of computer science focused on creating systems capable of performing tasks that would require intelligence if done by humans, such as problem-solving, decision-making, and language understanding.
	Algorithms	A well-defined set of step-by-step rules or instructions that a computer follows to process input data and produce an output or solution to a specific problem.
	Machine Learning (ML)	A branch of AI that enables systems to improve their performance on tasks by learning patterns and relationships from data rather than being explicitly programmed with fixed rules.

Category	Term	Definition
	Deep Learning (DL)	A specialized subset of machine learning that uses multi-layered artificial neural networks to automatically learn hierarchical representations of data, allowing for the detection of highly complex patterns.
	Natural Language Processing (NLP)	The field of AI that combines linguistics and computer science to enable machines to understand, interpret, and generate human language in a meaningful and useful way.
	Robotics	The integration of AI with physical machines designed to perceive their environment, process information, and take actions, enabling them to perform tasks autonomously or semi-autonomously.
Types of Machine Learning	**Supervised Learning**	A machine learning approach where models are trained using labeled datasets that pair inputs with correct outputs, enabling the system to make accurate predictions on new data.

Category	Term	Definition
	Unsupervised Learning	A machine learning approach that analyzes unlabeled datasets to discover hidden structures, patterns, or groupings within the data without explicit outcome labels.
	Reinforcement Learning	A machine learning paradigm where an agent learns by interacting with its environment, receiving feedback in the form of rewards or penalties, and optimizing actions to maximize long-term success.
Types of AI (by Behavior)	**Reactive Machines**	The simplest form of AI, designed to respond to immediate inputs or stimuli without memory of past experiences or the ability to adapt over time.
	Limited Memory	AI systems that can make decisions by learning from historical data and recent experiences, but with memory that is temporary and not retained for long-term use.

Category	Term	Definition
Types of AI (by Technology)	**Theory of Mind (ToM)**	A projected stage of AI development in which machines would be able to recognize, understand, and respond to human emotions, beliefs, and intentions, enabling more sophisticated social interactions.
	Self-Aware AI	A hypothetical stage of AI in which systems would achieve consciousness, self-recognition, and awareness of their internal states, allowing for autonomous decision-making at a human-like or higher level.
	Artificial Narrow Intelligence (ANI)	AI systems designed and trained for a specific, limited task, demonstrating competence within that domain but unable to generalize beyond it.
	Artificial General Intelligence (AGI)	A theoretical form of AI that would possess human-like intelligence, enabling it to understand, learn, and apply knowledge flexibly across a wide variety of tasks and contexts.

Category	Term	Definition
	Artificial Superintelligence (ASI)	A speculative form of AI that would surpass human intelligence in all areas, including reasoning, creativity, problem-solving, and emotional intelligence, with profound societal implications.

Key Takeaways

History and Evolution of AI:

- AI research began with Alan Turing's question of machine intelligence
- Key milestones: Turing test (1950), AI term coined (1955), Logic Theorist development
- Evolution from rule-based systems to machine learning to deep learning
- Current focus on neural networks and complex pattern recognition

Understanding AI (What is AI?):

1. Technological Perspective: • Focus on computational problem-solving • Emphasis on data processing and goal-driven actions • Flexibility in handling uncertain information
2. Psychological Perspective: • Emulation of human cognitive tasks • Focus on perception, memory, reasoning, and learning • Applications in object recognition and decision-making
3. Cognitive vs. Behavioral Perspective: • Four approaches: Acting/Thinking Humanly/Rationally • Balance between human-like behavior and rational decision-making • Recognition of human cognitive limitations

How AI Works:

- Core focus on data analysis and pattern recognition
- Learning through algorithmic processing
- Use of labeled and unlabeled datasets
- Continuous adaptation and improvement through experience

Think Deeper:

1. How does the progression from reactive machines to self-aware AI reflect our understanding of human intelligence?
2. What are the implications of developing AI systems with emotional understanding capabilities?

3. How might the evolution from ANI to ASI impact society and human-machine interaction?

2.

ETHICS IN AI AND MARKETING

AI Ethics

This text has been adapted from Floridi et al (2021) and Hermann (2021). All images in this chapter comes from this article.

The discussion about the **moral and ethical implications** of AI has been ongoing since the 1960s. The rapid advancement of Artificial Intelligence (AI) has revolutionized the field of marketing, offering unprecedented opportunities for personalization, efficiency, and customer engagement. However, with these technological strides come significant ethical considerations that demand our attention. This chapter delves into the complex ethical landscape of AI in marketing, examining the intricate balance between innovation and responsibility. We will explore five fundamental ethical principles – **beneficence, non-maleficence, autonomy, justice, and explicability** – and how they apply to AI-driven marketing strategies. As we

navigate through the potential benefits and risks, we'll uncover the tensions between individual, corporate, and societal interests. From privacy concerns and environmental impacts to issues of bias and transparency, this chapter aims to provide a comprehensive understanding of the ethical challenges faced by marketers, developers, and policymakers in the age of AI. By critically examining these issues, we lay the groundwork for **responsible AI implementation** in marketing, ensuring that as we harness the power of this technology, we do so in a manner that respects ethical boundaries and promotes the greater good.

We explore the ethical implications and concerns of using AI in marketing from multiple perspectives, including those of **companies, customers, and society** at large, including **environmental impacts**. This approach allows us to evaluate the validity and applicability of ethical principles across different stakeholder levels and to identify any tensions that might arise due to differing stakeholder interests.

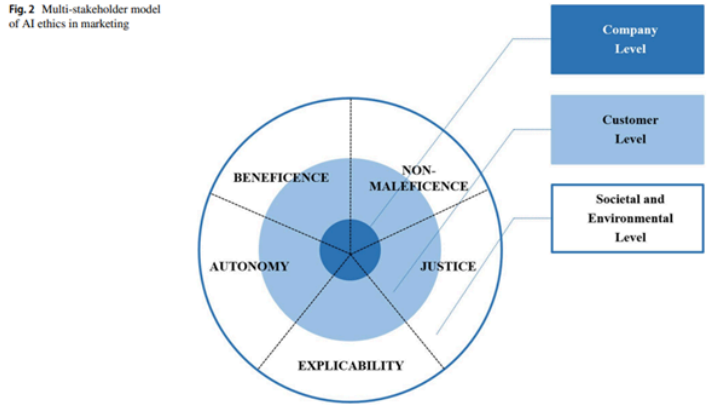

Fig. 2 Multi-stakeholder model of AI ethics in marketing

Floridi and colleagues have condensed ethical principles into five key areas: beneficence, non-maleficence, autonomy, justice, and explicability.

Beneficence involves promoting well-being and the common good, while **non-maleficence** focuses on preventing harm, whether accidental or intentional. These principles are not opposites but coexist to ensure AI benefits society while avoiding harm.

Autonomy concerns the balance between human and AI decision-making power, ensuring people can make choices freely without coercion.

Justice involves fairness, avoiding biases, and addressing past inequalities. It also relates to sharing AI's benefits and fostering solidarity.

Explicability means understanding how AI works and who is responsible for its actions. This principle supports and

enables the other ethical guidelines by ensuring AI systems are intelligible and accountable.

The need for explicability grows as AI systems become more complex and opaque. Understanding how AI functions and assigning responsibility is crucial for evaluating AI's benefits and risks, deciding on human and AI agency, and ensuring accountability for any failures or biases.

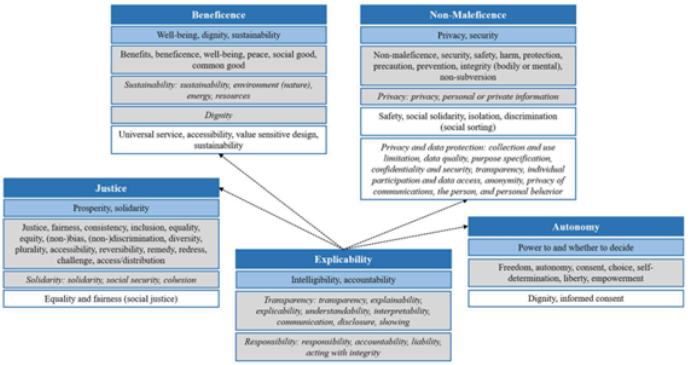

Fig. 1 AI ethics map. Principles in blue boxes are taken from Floridi et al. (2018), principles in gray boxes are taken from Jobin et al. (2019), and principles in white boxes are taken from Wright (2011). Principles in *italics* were not subsumed under *beneficence*, *non-maleficence*, *autonomy*, *justice*, or *explicability*, but listed as independent principles by Jobin et al. (2019) and Wright (2011), respectively

We now discuss each in depth.

Beneficence

The principle of beneficence—doing good and promoting well-being—is central to ethical considerations in AI-powered marketing. This principle raises important questions about who benefits from AI marketing technologies and whether these benefits can be reconciled across different levels of

impact. Let's examine how beneficence manifests in AI marketing and the complex ethical considerations it presents.



- The growing use of AI technologies like chatbots raises concerns about their benefice, as their increasing carbon emissions harm the environment

Benefits at the Individual Level

AI enables unprecedented **personalization** in marketing through sophisticated data analysis and prediction. Marketing teams can now analyze customers' digital footprints—including social media activity and smartphone data—to understand psychological traits and preferences. This capability allows companies to create highly personalized advertising, product recommendations, and marketing messages that resonate with individual customers.

Recommendation systems represent a prime example of

AI beneficence in marketing. These systems analyze customers' past behaviors and preferences, along with data from similar customers, to predict what products or services an individual might value. This benefits customers by filtering vast amounts of information and streamlining purchase decisions. For companies, these systems help convert browsers to buyers and facilitate cross-selling opportunities.

The Dual Nature of Benefits

The principle of beneficence in AI marketing manifests as a **two-way exchange of value**. Companies gain through increased sales efficiency and better customer targeting, while customers receive more relevant recommendations and a more personalized shopping experience. This mutual benefit suggests that AI marketing fulfills a fundamental aspect of beneficence: it serves a clear purpose and provides tangible advantages to both parties.

However, the concept of "good" in beneficence isn't always straightforward. For instance, when recommendation systems use patterns from similar customers to influence individual choices, they create a form of **indirect social influence**. If customers fully understood these underlying mechanisms, their perception of the benefits might change. The "black box" nature of many AI systems can obscure this understanding, raising questions about whether true beneficence requires **transparency**.

The Challenge of Competing Benefits

A critical ethical challenge emerges when considering beneficence across different levels of impact. What benefits individuals may not benefit society or the environment as a whole. AI marketing systems typically optimize for individual customer satisfaction and increased consumption, which can conflict with **broader societal goals**.

For example, while AI-powered marketing might benefit individual customers through personalized shopping experiences, it can simultaneously contribute to environmental challenges through increased consumption. E-commerce platforms using AI recommendations may drive sales efficiently, but this efficiency comes with environmental costs in terms of packaging waste, transportation emissions, and energy consumption.

The Environmental Dimension

The environmental impact of AI-driven consumption presents a particular challenge to the principle of beneficence. While meeting individual consumer needs provides clear benefits, the resulting resource depletion and environmental damage raise questions about broader societal good. Online retail, powered by AI recommendations, often has a **larger ecological footprint** than traditional shopping due to factors like individual packaging and last-mile delivery.

Additionally, the technology itself consumes significant

energy in its development and operation. These environmental costs must be considered when evaluating whether AI marketing truly fulfills the principle of beneficence on a societal level.

Balancing Individual and Collective Benefits

The tension between individual and collective benefits creates an ethical dilemma in AI marketing. While AI systems excel at maximizing **individual utility** through targeted recommendations based on past preferences (exploitation), they might better serve society by exploring and promoting more **sustainable alternatives** (exploration). This raises the question of whether true beneficence should prioritize immediate individual benefits or longer-term societal good.

Future Considerations

Understanding beneficence in AI marketing requires acknowledging these multiple layers of impact. While the technology clearly provides benefits at the individual level, its broader implications for society and the environment complicate the ethical equation. Moving forward, the challenge lies in developing AI marketing systems that can balance individual benefits with collective good, perhaps by incorporating **sustainability** and **social responsibility** into their optimization criteria.

The principle of beneficence thus demands a holistic view that considers both immediate individual benefits and broader societal impacts. This understanding can help guide the development of more ethical AI marketing practices that truly serve the greater good while maintaining their value to individuals and businesses.

Non-Maleficence

The principle of non-maleficence—avoiding harm—presents critical considerations in AI-powered marketing. While AI marketing systems can offer significant benefits, they also carry potential risks that must be carefully managed. Understanding these risks helps organizations implement AI marketing technologies more responsibly.



PETER NORTON TRANSPORTATION SEP 13, 2018 7:00 AM

Self-Driving Car Developers Should Put Pedestrians First

Opinion: Contrary to the hype, walkers and bikers are much smarter than autonomous vehicles

NATALIE BEHRING/REUTERS

- This article highlights the principle of non-maleficence, showing how prioritizing pedestrian safety in self-driving cars reflects the need to manage risks responsibly in AI-powered technologies

The Relationship with Beneficence

Non-maleficence and beneficence are closely linked when considering AI marketing's broader impacts. From an environmental and societal perspective, these principles often align—when AI-driven marketing fails to promote environmental good, it simultaneously risks causing

environmental harm. However, at the individual and company level, the relationship becomes more complex. AI applications can simultaneously benefit users while potentially causing other forms of harm, creating an **ethical tension** that requires careful navigation.

Privacy and Data Protection Concerns

A primary consideration in non-maleficence is the protection of personal privacy and data security. Privacy risks in AI marketing can emerge in three key ways:

First, during data collection, companies might gather information without proper informed consent from customers. Second, stored data might be compromised through breaches or de-anonymization. Third, AI systems can draw potentially invasive inferences both directly from individual customer data and indirectly through interaction data with other customers, particularly in collaborative filtering systems.

The **scale of data collection** in AI marketing amplifies these privacy concerns. While more data generally leads to better predictions and personalization, it also increases privacy risks. This creates a fundamental tension between marketing effectiveness and privacy protection. **Regulatory frameworks**, such as the European Union's General Data Protection Regulation (GDPR), attempt to address these concerns by requiring proactive privacy protection measures in the design and implementation of AI systems.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=193#oembed-1

- The Target pregnancy scandal serves as a stark example of how companies have used machine learning technologies in marketing to invasively predict personal information, raising significant concerns about consumer privacy and ethical data usage.

Accuracy and Data Quality Issues

The non-maleficence principle also encompasses the importance of accuracy and data quality in AI marketing systems. These systems can only be as reliable as their underlying data. When input data contains biases, inaccuracies, or errors, it can lead to flawed conclusions and potentially harmful decisions.

A particular concern arises when marketing decisions are based on correlations found in large datasets without establishing causality. Making individual-level decisions based

on population-level patterns can be problematic and potentially harmful to customers.

The Challenge of Algorithm Reliance

The relationship between customers and algorithmic recommendations presents another area where the principle of non-maleficence comes into play. Two key challenges emerge:

- **Algorithm Overreliance**: When customers depend too heavily on AI recommendations, they might accept inferior suggestions that diminish their well-being. This represents a "false positive" scenario where customers trust and follow recommendations that don't truly serve their interests.
- **Algorithm Aversion**: Conversely, customers might reject valid AI recommendations due to skepticism or distrust, potentially missing out on beneficial opportunities. This "false negative" scenario can impact both customer well-being and company objectives.

Pricing and Competition Concerns

AI systems used in pricing decisions can also raise non-maleficence concerns. Without proper oversight, pricing algorithms might engage in tacit collusion or overestimate price sensitivities, leading to **artificially inflated prices**. This can harm not only individual customers but potentially

contribute to broader economic issues through increased price levels across markets.

The Black Box Challenge

The complexity of AI systems—often referred to as the "black box" problem—makes it particularly difficult to assess potential harms. Both customers and companies may struggle to understand the extent of data capture, evaluate data quality, or assess model specifications. This lack of transparency complicates efforts to ensure non-maleficence in AI marketing applications.

Moving Forward

To uphold the principle of non-maleficence, organizations implementing AI marketing systems must:

- Implement robust privacy protection measures
- Ensure high data quality standards
- Carefully monitor algorithmic decisions
- Maintain transparency where possible
- Regularly assess potential negative impacts

Understanding and addressing these concerns helps organizations balance the benefits of AI marketing with their ethical obligation to avoid harm. This balanced approach is

essential for maintaining customer trust and ensuring the sustainable development of AI marketing technologies.

Autonomy

The principle of autonomy in AI marketing centers on preserving **human agency** and decision-making power. As AI systems become increasingly sophisticated in their ability to influence and shape consumer choices, understanding how to maintain meaningful human autonomy while leveraging AI's capabilities presents a crucial challenge in modern marketing.



- Algorithms, such as those used on YouTube, influence over 90% of watch time, raising concerns about whether

users truly have control over their choices or if AI systems are undermining human decision-making power

The Foundation of Consumer Autonomy

Consumer autonomy fundamentally refers to individuals' ability to make and execute decisions **independently**, without undue external influence. In the context of AI ethics, this transforms into a meta-autonomy framework where humans retain the power to choose when to make decisions themselves and when to delegate decision-making to AI systems. This ability to choose when to delegate represents a crucial aspect of preserving human agency in an AI-enabled world.

Corporate Governance and Human Oversight

From an organizational perspective, maintaining human autonomy requires robust governance mechanisms that keep humans actively involved in AI systems' operation. This becomes particularly critical when AI operates in **ethically sensitive contexts**, such as when systems engage in emotional interactions with customers or face moral decisions. While AI continues to advance in emotional intelligence and moral reasoning capabilities, current systems primarily recognize and simulate emotions rather than truly understand them. Until AI develops genuine emotional intelligence and moral agency, **human oversight** remains essential.

The Influence on Consumer Decision-Making

AI marketing systems can significantly influence consumer decision-making through several key mechanisms. Through **personalization** and **targeting**, these systems shape the information and options presented to consumers, effectively altering the choice architecture within which decisions are made. Recommendation systems act as information filters, pre-selecting options based on predicted preferences and past behaviors. By determining what information and options consumers see, AI systems can function as "nudges" that influence decision-making in subtle but powerful ways.

The Delegation Challenge

When consumers interact with AI marketing systems, they often delegate part of their decision-making power at the information-gathering stage. This delegation can offer benefits through increased efficiency and personalization, saving time and cognitive resources. However, it also presents potential risks. Consumers might become too dependent on AI recommendations, potentially accepting suboptimal choices without sufficient consideration. Conversely, some consumers might reject helpful AI guidance due to distrust or skepticism. There's also potential for AI systems to present manipulated or deceptive content that unduly influences consumer choices.

Preserving Meaningful Choice

Organizations must carefully design their AI marketing systems to preserve genuine **consumer autonomy**. This includes providing clear options about levels of AI engagement, maintaining transparency about how AI systems influence the decision-making environment, and ensuring consumers can opt out of AI-driven interactions when desired. A crucial element in preserving autonomy is ensuring that both consumers and company representatives can make **informed decisions** about when to delegate to AI systems. This requires a basic understanding of how these systems function – their capabilities, limitations, and potential biases.

Looking Forward

As AI systems become more sophisticated in their ability to understand and respond to human emotions, maintaining meaningful human autonomy will become increasingly complex. The key lies in striking a balance between leveraging AI's capabilities to improve marketing effectiveness while preserving consumers' fundamental right to make independent, informed decisions. This balance requires ongoing attention to transparency, choice architecture, and human oversight in AI marketing systems.

Organizations implementing AI marketing must recognize that true autonomy extends beyond simply offering choices – it requires ensuring those choices are **meaningful and**

informed. This means maintaining clear processes for human intervention and decision-making, especially in sensitive customer interactions. As AI technology continues to evolve, the challenge of preserving human autonomy while harnessing AI's benefits will remain at the forefront of ethical marketing considerations.

Justice

The principle of justice in AI marketing addresses how these systems can either promote or undermine **fairness and equality**. As AI systems become increasingly prevalent in marketing decisions, understanding their potential to create or amplify discrimination becomes crucial for ethical implementation.

AI For Justice: How A Female Founder Builds The Case For Lawyers

Geri Stengel Contributor ⓘ
Geri Stengel writes about the success factors of women entrepreneurs.

Follow

0

Updated Mar 13, 2024, 07:00am EDT



Evgeniia Ben Artzi, cofounder and CEO, and Gila Hayat, cofounder and CTO at Datoma.ai

< FEATURES

March 21, 2023 1:33 pm



Racial bias in facial recognition algorithms

The Challenge of Algorithmic Bias

AI systems can inherit and amplify human biases, presenting a significant ethical challenge in marketing. These biases often manifest in personalization strategies, psychological targeting, and customer relationship management. The key concern is that AI systems may discriminate against certain customer groups based on demographic, psychological, or economic factors.

Several factors can contribute to discriminatory outcomes in AI marketing systems. When certain demographic groups are over- or underrepresented in training data, AI systems may **develop biased predictions**. This can occur when data collection is skewed or when certain groups are less willing to share data due to privacy concerns. Additionally, initial biases in data can create a cycle where biased predictions inform decisions, which then generate new biased data, perpetuating and potentially amplifying the original bias. Customer prioritization based on income or profitability can **exacerbate existing social inequalities** by limiting certain groups' access to opportunities or information.

Impact Across the Marketing Ecosystem

The impact of AI bias extends throughout the marketing ecosystem, affecting various stakeholders differently. For consumers, AI marketing systems can create or reinforce disparities based on gender, age, race, and economic status.

This manifests through **differential targeting, pricing, or access** to products and services. Particularly concerning is the potential targeting of vulnerable consumer groups, which could exploit existing disadvantages.

On the corporate level, AI systems like recommendation engines can create **market disparities**. They might reduce product diversity by focusing consumer attention on popular items, potentially harming smaller businesses or those offering niche products. Additionally, businesses that cannot effectively implement AI marketing systems may face competitive disadvantages. The rise of AI-powered e-commerce platforms can create market access barriers for traditional retailers, potentially leading to market concentration and reduced competition.

The Challenge of Transparency

The "black box" nature of many AI systems makes it difficult to identify and address biases. This lack of transparency creates additional challenges in ensuring justice, as it becomes harder to detect when systems are producing discriminatory outcomes. Organizations must work to implement rigorous protocols for data collection and validation to ensure **representative samples** across all demographic groups. Regular monitoring of AI systems for potential discriminatory outcomes and establishing processes for **correcting identified biases** becomes crucial.

Promoting Fairness in AI Marketing

While much attention focuses on preventing discrimination, AI can also serve as a tool for promoting justice. Recent research has demonstrated that algorithms can be used to **detect discrimination** and help de-bias human judgments, suggesting potential positive applications for promoting fairness. Organizations can incorporate ethics specialists into AI development teams to help identify and address potential justice concerns early in the development process.

Looking Forward

Success in AI marketing should be measured not only by business metrics but also by the system's ability to serve all stakeholders fairly and ethically. Organizations must remain vigilant in monitoring and addressing potential discriminatory outcomes while exploring ways to use AI as a tool for promoting greater fairness in marketing practices. This requires balancing **efficiency and fairness** by considering both business objectives and ethical implications when designing AI marketing strategies.

Regular assessment of how AI decisions affect different customer segments and stakeholder groups becomes essential, as does maintaining human supervision of AI systems, particularly in decisions that could affect vulnerable groups. The justice principle reminds us that while AI marketing systems can enhance efficiency and personalization, they must

be designed and implemented with careful attention to fairness and equality.

In conclusion, organizations must commit to ethical oversight, transparent practices, and regular assessment of AI systems' impact on various stakeholder groups. Only through this ongoing commitment can AI marketing systems truly serve their purpose while maintaining fairness and equality across all stakeholder groups.

Explicability

Explicability stands as one of the most challenging and debated principles in AI marketing ethics. This principle addresses how AI systems can be made understandable and accountable to stakeholders, particularly given their often opaque "black box" nature. As AI systems increasingly make important decisions affecting customers and businesses, the need for explicability becomes more crucial than ever.



The Challenge of Black Box AI

AI marketing systems often operate as black boxes, where the relationship between inputs and outputs remains unclear to both customers and businesses. This opacity raises significant concerns, particularly when these systems handle sensitive personal data or make high-stakes decisions about customer targeting, pricing, or recommendations. The complexity of modern AI systems can make it difficult for even their developers to fully understand how certain decisions are reached, creating challenges for accountability and trust.

Key Dimensions of Explicability

Explicability encompasses two fundamental dimensions that work together to create transparency and trust in AI marketing systems. The first is **intelligibility**, which refers to helping

humans understand how AI systems function without necessarily diving into their internal complexities. Intelligibility serves as an enabling factor for other ethical principles, as understanding how AI works helps people evaluate its benefits, potential harms, fairness, and appropriate use.

The second dimension is **accountability**, which addresses who bears responsibility for AI systems' decisions and outcomes. As AI systems increasingly shape marketing decisions, establishing clear lines of accountability becomes crucial, especially when outcomes prove adverse or discriminatory. Organizations must be prepared to take responsibility for their AI systems' actions and provide recourse when things go wrong.

Balancing Competing Interests

Organizations implementing explicability face several significant challenges. Technical complexity makes explaining sophisticated AI systems in simple terms difficult without losing accuracy or meaningful detail. While customers deserve explanations about how AI affects them, too much technical detail can lead to confusion and frustration rather than understanding. Additionally, complete transparency might conflict with protecting proprietary information and maintaining competitive advantages.

Some AI applications, such as chatbots, might become less effective if their artificial nature is made too explicit, creating

a tension between transparency and performance. This raises important questions about how to balance the ethical imperative for transparency with practical business considerations.

Practical Implementation in Marketing

Implementing explicability in AI marketing requires a thoughtful approach that considers various stakeholder needs. Organizations should provide simple, clear explanations of how their AI systems work, focusing on what customers need to know rather than technical details. This might include explaining why certain recommendations are made or how personal data influences marketing decisions.

Establishing clear accountability structures for AI system outcomes is equally important. Organizations must be prepared to take responsibility when AI systems make mistakes or produce undesired outcomes. This includes having processes in place for addressing complaints and correcting errors.

The Impact on Other Ethical Principles

Explicability's importance extends beyond mere transparency, as it enables the evaluation of other ethical principles in AI marketing. Understanding how AI systems work helps assess their benefits and potential harms, supporting the principles of beneficence and non-maleficence. Transparency about AI

operations helps identify potential biases and discriminatory outcomes, supporting justice. Clear explanations enable informed decisions about whether to delegate choices to AI systems, supporting autonomy.

The Future of Explicable AI

Some experts advocate for building inherently interpretable AI systems rather than trying to explain complex black box systems after the fact. This approach suggests designing AI marketing systems with explicability in mind from the start, rather than attempting to add transparency later. This might mean choosing simpler, more interpretable algorithms over more complex ones, even if they offer slightly lower performance.

Looking forward, organizations must continue to develop new approaches to explicability that balance stakeholder needs with practical constraints. This might include innovative ways of visualizing AI decision-making processes or new frameworks for establishing accountability. Success in AI marketing increasingly depends on finding the right balance between complexity and comprehensibility, between transparency and effectiveness, and between disclosure and competitive advantage.

In conclusion, explicability remains a crucial principle for ethical AI marketing, even as implementing it presents significant challenges. Organizations that successfully navigate these challenges while maintaining **transparency and**

accountability will likely find themselves better positioned to build trust with customers and other stakeholders in an increasingly AI-driven marketing landscape.

Interdependencies Between Ethical Principles and Levels of Intelligence

This conceptual analysis reveals that ethical principles related to AI in marketing interact and sometimes conflict, so they cannot be judged in isolation but in relation to each other.

Explicability is an enabling principle for beneficence, non-maleficence, justice, and autonomy. At the same time, justice and autonomy shape how beneficence and non-maleficence are applied. AI systems can be both beneficial and harmful, depending on the stakeholders involved. For example, AI-based personalization strategies might satisfy customer needs and boost company sales (beneficence), but they also raise the overall consumption level, harming the environment and society (non-maleficence). Privacy and data protection issues can compromise the benefits of personalization for both customers and companies.

The complexity of ethical issues varies across AI applications, technological sophistication, and pervasiveness. Ethical challenges and tensions may increase with the level of AI intelligence and humanization, transitioning from

cognitive (analytical, mechanical) to emotional (feeling) and social intelligence. This progression can lead to greater benefits, such as better-targeted customer needs, but also increases the issues related to explicability, such as accountability in case of failures.

Ethical challenges related to justice and autonomy can also increase with more sophisticated AI, though this is not always the case. For example, AI can be used as a discrimination detector, identifying and mitigating biases. However, the need for human agency and oversight becomes more critical, especially in ethically significant contexts. Achieving non-maleficence will depend on how customer data is gathered and treated.

As AI systems become more human-like, they may be perceived with anthropomorphism and mind perception, attributing experience, warmth, agency, and competence to AI. Intense use and interaction with human-like AI can lead to psychological ownership and emotional attachment, which might have detrimental effects if AI deployment gets out of control or if there is overreliance on AI systems. Extensive use can also lead to the dehumanization of users.

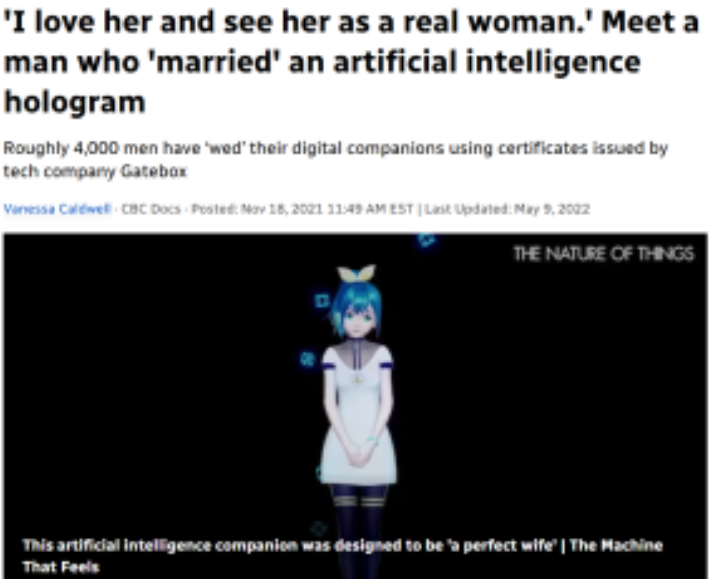

- The phenomena of humans falling in love with AI companions highlights the potential societal and individual implications of overreliance on human-like AI systems

From a deontological perspective, AI applications might be questioned if non-maleficence cannot be guaranteed, which could lead to missing out on opportunities to serve customers and achieve beneficence at the customer and company levels. A strictly deontological approach may not adequately address the complex ethical concerns involved. Instead, a utilitarian perspective, which weighs benefits and costs across all stakeholders, might be more appropriate. This approach can

consider multiple values, objectives, and utilities at individual, group, and societal levels.

AI systems could incorporate multi-objective maximum-expected-utility concepts aligned with human values and ethical principles. However, this technological design and implementation are challenging. Determining which utilities to optimize is difficult and varies across stakeholders. Moreover, artificial moral agency is still in its early stages, requiring significant human oversight. Humans must decide how to align AI applications and systems with ethical principles and human values and determine whether AI systems should base their ethical decision-making on pre-defined ethical theories (top-down), flexible self-learning mechanisms based on certain values (bottom-up), or a combination of both (hybrid).

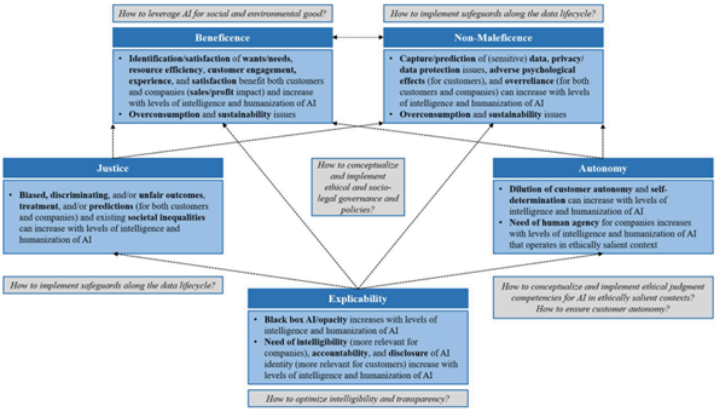

Fig. 3 Interdependencies of ethical principles and levels of intelligence and future research questions

AI for Societal Good?

AI in marketing creates a challenging balance between effective selling and environmental responsibility. When AI personalizes marketing messages and manages customer relationships well, it often leads to increased consumption. While this benefits businesses, it conflicts with the growing need for sustainable consumption patterns and the call for people to "buy better but less."

Marketing and sustainability might seem to be working against each other at first glance. After all, marketing typically aims to increase sales, while sustainability often requires reducing consumption. However, AI offers promising ways to align these seemingly opposing goals.

Consider how AI can support more sustainable consumer choices. Through advanced data analysis, AI can identify customers who are interested in sustainable products and match them with appropriate eco-friendly offerings. The technology can process complex relationships between what companies offer (like product features and pricing) and what influences consumer decisions (such as personal values and demographics).

However, the relationship between marketing and sustainable consumption is complex. Research shows interesting patterns in how consumers respond to sustainable products. For instance, when people care about others as much as themselves, they often value sustainable products

differently. Surprisingly, when environmentally friendly products are put on sale, some consumers become less likely to buy them. This might be because discounts can make these products seem less premium or authentic.

AI can help marketers navigate these complexities. For example, it can analyze when sustainability messages are most effective in a customer's buying journey. While some customers might respond well to environmental messaging when first discovering products, others might be more influenced by it when making their final purchase decision.

The challenge lies in overcoming psychological barriers that prevent sustainable behavior change. Even when people want to make environmentally conscious choices, various mental blocks can get in the way. AI-powered marketing can help identify and address these barriers, making it easier for consumers to align their purchasing decisions with their environmental values.

One promising way AI can promote sustainability is through **psychological targeting**. This means using AI to identify consumers who are most likely to be interested in sustainable and ethical products based on their psychological characteristics. By analyzing digital behavior patterns, AI can help marketers customize their messages to resonate with environmentally conscious consumers.

AI can even help set appropriate prices for sustainable products. By analyzing consumers' digital footprints, marketers can better understand how much different

customer groups are willing to pay for environmentally friendly options. This is particularly important because research shows that pricing strategies can significantly impact how consumers respond to eco-friendly products.

However, with this power comes significant responsibility. While psychological targeting can promote positive social and environmental outcomes (beneficence), marketers must be careful to avoid potential harm (non-maleficence). For example, it would be unethical to use these targeting capabilities to exploit vulnerable individuals or those prone to compulsive buying behaviors.

The challenge of ethical targeting becomes even more complex when considering AI systems themselves. These systems can unintentionally discriminate against certain groups if they're trained on biased data or if their mathematical models aren't properly designed. To prevent unfair outcomes that could undermine the positive purpose of these AI applications, marketers need to understand both the data being used and how their AI systems work.

This creates a shared responsibility between marketers and technical teams. Marketers, who are often held accountable for the outcomes, need to be aware of potential biases in their data and AI systems. Meanwhile, data scientists and AI developers must use their technical expertise to identify and fix any biases or errors in the systems they create.

Think of it like a powerful tool that needs careful handling: Just as a sharp knife can be used to prepare a healthy meal

or cause harm, AI-powered marketing can either promote sustainable consumption or exploit vulnerable consumers. The key is understanding the tool and using it responsibly.

AI-powered **recommendation systems** can act as digital guides toward sustainable consumption. Just like personalized marketing messages, these systems can nudge consumers toward more environmentally conscious choices. The key is balancing two approaches: helping customers explore new sustainable options while respecting their existing preferences.

For example, when showing product recommendations, an AI system could introduce sustainable alternatives alongside conventional products. This approach gradually expands the variety of environmentally friendly options presented to consumers, whether they're shopping for fast fashion or considering more sustainable, longer-lasting clothing choices.

AI systems can also help prevent compulsive buying behaviors by monitoring purchase patterns. When the system detects potential issues, such as unusually frequent purchases or repeated buying of similar items, it can provide informational nudges. These gentle reminders don't force behavior change but instead **encourage thoughtful decision-making** by highlighting the environmental or personal impact of excessive consumption. Research shows that simply prompting consumers to think about their existing possessions can reduce impulse buying.

When designing these AI systems, companies must carefully

balance effectiveness with consumer autonomy. There are two main approaches to nudging:

1. **Informational nudges:** Changing the type of information shown to consumers
2. **Structural nudges:** Modifying how choices are presented

While structural nudges tend to be more effective, informational nudges better preserve consumer autonomy by allowing individuals to make their own informed decisions.

This raises an important question in AI development: How should we build ethics into these systems? Two main approaches exist:

- **Ethics by design:** Building ethical constraints directly into the system's structure
- **Pro-ethical design:** Providing information that enables ethical decision-making

The choice between these approaches depends on several factors, including the need for transparency, accountability, and consumer acceptance. Success ultimately requires collaboration between AI developers and ethics experts to translate ethical principles into practical applications that benefit both consumers and the environment.

Key Terms

Category	Term	Definition
Ethical Principles	**Beneficence**	Using AI to deliver clear benefits for individuals, organizations, and society.
	Non-maleficence	Preventing avoidable harm from AI, including privacy loss, manipulation, unsafe behavior, and environmental damage.
	Autonomy	Preserving people's ability to choose, understand options, and say no to AI recommendations.
	Justice	Ensuring fair treatment and equal access while avoiding discrimination across groups.
	Explicability	Making AI understandable and accountable, with reasons for decisions and clear responsibility.
Privacy & Data	**Privacy**	People's control over how their personal data is collected, used, and shared.

Category	Term	Definition
	Informed Consent	Freely given permission to process data after clear explanation of purpose, risks, and choices.
	Data Minimization	Collect only what is needed and retain it only as long as necessary.
	Re-identification	Linking anonymous records back to real people, which undermines privacy.
Fairness & Governance	**Algorithmic Bias**	Unfair outcomes caused by skewed data, labels, features, or objectives.
	Black box model	A model whose reasoning is hard to interpret, making explanations and audits difficult.
	Accountability	Clear responsibility for AI outcomes and a path to fix problems.
	Human-in-the-loop	Humans supervise and can approve, correct, or override AI outputs.
Autonomy & Influence	**Choice Architecture**	How options, defaults, and information are presented, which shapes decisions.

Category	Term	Definition
	Nudge	A small change in presentation that steers choices without removing options.
	Delegation	Deciding which steps are handled by AI and which stay with people.
Sustainability	**Sustainable Consumption**	Buying choices that reduce environmental harm across a product's life cycle.
	AI Carbon Footprint	Energy use and emissions from training and running AI systems.

Summary: Ethics in AI and Marketing

Key Takeaways:

– Five fundamental ethical principles in AI marketing: beneficence, non-maleficence, autonomy, justice, and

explicability.

– AI in marketing offers benefits like personalization but raises concerns about privacy, environmental impact, and bias.

– Ethical considerations vary across stakeholder levels (individual, corporate, societal) and AI sophistication.

– Transparency and accountability are crucial but challenging as AI systems become more complex.

– A balanced approach considering multiple perspectives is necessary for ethical AI implementation in marketing.

Connections:

This chapter builds on the AI concepts introduced in previous chapters, exploring their ethical implications in marketing contexts. It sets the foundation for responsible AI use in marketing strategies discussed in later chapters.

AI in Action:

Recommender systems used by e-commerce platforms demonstrate both the benefits (personalized recommendations) and potential drawbacks (privacy concerns, environmental impact) of AI in marketing.

Think Deeper:

1. How can marketers balance the benefits of AI-driven personalization with concerns about consumer autonomy and privacy?

2. Consider the environmental impact of AI-driven

marketing. How might marketers address these sustainability concerns while still leveraging AI's benefits?

Further Exploration:

– Research real-world examples of companies implementing ethical AI practices in their marketing strategies.

– Explore emerging regulations and guidelines for ethical AI use in marketing across different countries or regions.

3.

AI AND THE TRANSFORMATION OF CONSUMER EXPERIENCES

This chapter is based on Puntoni et al. (2021)

As artificial intelligence becomes increasingly woven into the fabric of daily life, a crucial shift is occurring in how we must understand its impact. While much attention has focused on AI's technical capabilities and underlying algorithms, the more pressing question now centers on how consumers actually **experience** and **interact** with these technologies. This experiential lens reveals that AI's influence extends far beyond mere utility, touching on fundamental aspects of human identity, autonomy, and social connection.

Consider how AI shapes our daily routines: We wake up to personalized news feeds, navigate our commutes with AI-optimized routes, work alongside AI-powered tools, and end our days streaming content selected by recommendation algorithms. Each of these interactions represents more than

just a technological convenience—it embodies a **complex relationship** between human and machine that can simultaneously empower and constrain, connect and alienate, serve and exploit.

This chapter examines four fundamental dimensions of the consumer AI experience: **data capture, classification, delegation, and social interaction**. Each dimension creates distinct tensions in how consumers engage with AI systems. When sharing personal data, consumers face the paradox of being both served and exploited. As AI systems classify and categorize them, they wrestle with feeling understood versus misunderstood. In delegating tasks to AI, they balance the promise of empowerment against the fear of replacement. And through social interactions with AI, they navigate new forms of connection that can either enhance or diminish human relationships.

Understanding these tensions is crucial for both practitioners and researchers. For businesses, success in the AI era requires moving beyond technical performance metrics to consider how their systems affect consumers' sense of **identity, agency, and well-being**. For scholars, these experiential dimensions offer rich territory for exploring how AI is reshaping fundamental concepts of **consumer behavior** and **human-computer interaction**.

Through both **sociological** and **psychological** lenses, this chapter unpacks each of these core tensions, examining how they manifest in consumer experiences and what they mean

for the future of AI development. By understanding these dynamics, we can work toward AI implementations that enhance rather than diminish human capability and dignity, while creating value for both consumers and organizations.

Consumer Experience and AI Capabilities

To better understand the **consumer AI experience**, it is critical to shift focus from the underlying technology to how consumers engage with these capabilities. Consumer experience refers to the interactions between consumers and companies during the customer journey, encompassing emotional, cognitive, behavioral, sensorial, and social dimensions (Brakus, Schmitt, and Zarantonello 2009; Lemon and Verhoef 2016). This perspective broadens our understanding of AI's role in consumer lives beyond utility and function to include its emotional and symbolic significance (Mick and Fournier 1998).

A framework proposed by researchers outlines four experiences that reflect consumer interactions with AI capabilities:

1. **Data Capture** – The experience of sharing personal data with AI systems.
2. **Classification** – The experience of receiving predictions

or recommendations based on personal data.

3. **Delegation** – The experience of allowing AI to perform tasks on behalf of the consumer.
4. **Social Interaction** – The experience of communicating with AI systems as interactive partners.

The four AI capabilities – data capture, classification, delegation, and social interaction – each create distinct consumer experiences that are characterized by fundamental tensions. These tensions arise as AI simultaneously offers benefits while potentially undermining aspects of consumer autonomy and wellbeing:

- **Data Capture**: Tension between being served and exploited When AI systems collect consumer data, they can provide highly personalized services and experiences (served), but may also leave consumers feeling that their privacy has been invaded and their personal information commodified (exploited).
- **Classification**: Tension between being understood and misunderstood AI's ability to analyze and categorize consumers can lead to helpful personalization that makes them feel truly understood, but can also result in reductive or biased categorizations that leave consumers feeling mischaracterized or discriminated against.
- **Delegation**: Tension between being empowered and replaced While AI can empower consumers by handling

tedious or complex tasks, it can also create anxiety about human capabilities becoming obsolete or irrelevant.

- **Social Experience**: Tension between being connected and alienated AI-enabled social interactions can foster new forms of connection and engagement, but may simultaneously create experiences of exclusion or disconnection, particularly for marginalized groups.

Understanding these tensions through both sociological and psychological lenses is crucial for several reasons:

- The **sociological perspective** reveals how these tensions reflect and reshape broader **social structures and power dynamics**. For instance, data capture practices aren't just about individual privacy concerns, but about fundamental shifts in how personal information becomes a form of capital in what scholars call "surveillance capitalism."
- The **psychological perspective** illuminates how these tensions manifest in **individual consumer experiences and behaviors**. This helps explain why consumers might simultaneously embrace and resist AI technologies, or why the same AI capability might be experienced very differently by different individuals.

Each section of this chapter examines one of these core tensions through both lenses, followed by managerial

recommendations for addressing them. This dual analysis helps future marketers understand not only how to implement AI effectively, but how to do so in ways that benefit both individual consumers and society as a whole.

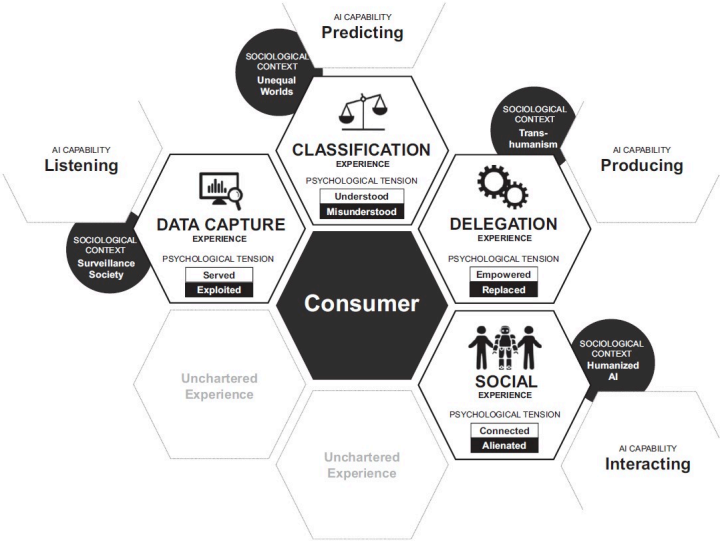

From Puntoni et al. (2021)

The AI Data Capture Experience

The **listening capability** in AI enables systems to gather data about consumers and their environments. This process forms the basis of the **data capture experience**, which involves various ways data is transferred to AI systems. Consumers may **intentionally** provide data, but their understanding of how

the data will be used often varies. Some share data with clarity and awareness of its purpose, while others surrender it amid high levels of uncertainty (Walker 2016). AI systems may also acquire data **passively**, gathering "digital shadows" that consumers leave behind in daily activities. For example, shoppers in stores equipped with facial recognition technology or users of iRobot Roomba vacuums—devices that map residential spaces—unknowingly contribute to data collection (Kuniavsky 2010).

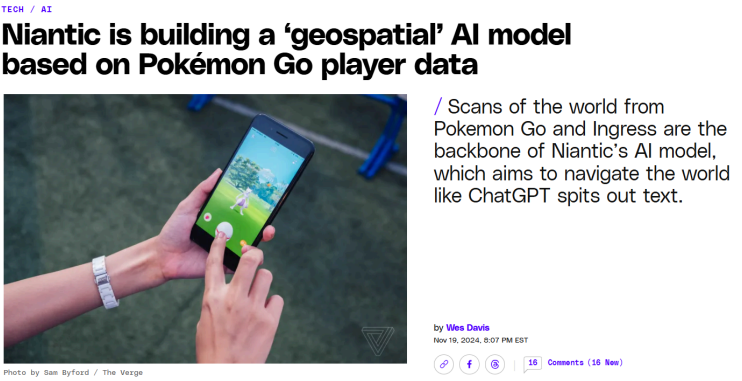

TECH / AI

Niantic is building a 'geospatial' AI model based on Pokémon Go player data

/ Scans of the world from Pokemon Go and Ingress are the backbone of Niantic's AI model, which aims to navigate the world like ChatGPT spits out text.

by Wes Davis
Nov 19, 2024, 8:07 PM EST

16 Comments (16 New)

Photo by Sam Byford / The Verge

Benefits of Data Capture

Despite these concerns, the data capture experience offers tangible benefits to consumers:

- **Access to Customized Services**: Sharing personal data enables tailored services, information, and entertainment, often provided for free. For example,

Google Photos captures users' memories and, in return, provides an AI-powered assistant that suggests context-specific actions, such as creating photo albums or sharing images.

- **Simplified Decision-Making**: AI systems can reduce the cognitive and emotional burden of decision-making by matching personal preferences with optimal choices. For instance, digital assistants can provide curated recommendations, sparing consumers the effort of evaluating numerous options (André et al. 2018).
- **Opportunities for Self-Improvement**: Customized AI-driven solutions can support personal development. One project from Alphabet integrates data from smartphones, genomes, wearables, and ambient sensors to deliver personalized healthcare recommendations (Kuchler 2020).

Challenges in Data Capture

The data capture process raises several concerns for consumers, often tied to the complexity and opacity of AI systems:

1. **Intrusive Data Collection**: Modern data acquisition methods are increasingly pervasive and difficult to avoid.
2. **Aggregation Over Time**: Even when consumers intentionally share data, they often remain unaware of how it is combined across contexts to create

comprehensive profiles.

3. **Transparency and Regulation**: Data brokers often operate in an unregulated environment, with limited accountability and little clarity about how data is used (Grafanaki 2017).

These factors can undermine consumers' sense of ownership over their personal data and erode their **personal control**—the feeling that individuals, rather than external forces, dictate outcomes (DeCharms 1968). This loss of control has both sociological and psychological implications, which are critical to understanding the consumer experience.

The Double-Edged Sword of Data Capture

While the benefits are significant, helping provide a **better service,** the data capture experience often leaves **consumers feeling exploited**. They may lack awareness of how their data is being used or feel powerless to control the process. Striking a balance between empowering consumers with customized services and addressing concerns about privacy and control is essential for businesses leveraging AI. By fostering transparency and ensuring ethical data practices, companies can mitigate the risks associated with data capture and build trust with consumers.

Sociological Context: The Surveillance Society Narrative

The issue of losing ownership over **personal data** is often linked to a perceived loss of personal control, fueled by concerns about technology's capacity to monitor and influence human behavior. This narrative has deep roots in popular culture, with dystopian stories like George Orwell's *1984* and Philip K. Dick's *Minority Report* portraying societies where constant surveillance **erodes privacy and individual autonomy**. Such depictions resonate with sociological critiques, which view data capture as central to the rise of a **surveillance economy** where personal information becomes a key form of capital (Zuboff 2019).

The Shift to a Surveillance Marketplace

A turning point in this narrative came when companies like Google began to reframe consumer data as a valuable economic resource. In the early 2000s, Google pioneered a business model that transformed consumer data—once a mere by-product of online activity—into an economic asset driving a new type of commerce. This commerce thrives on **data surplus**, which is processed using machine intelligence to create "prediction products" designed to anticipate consumer behaviors (Zuboff 2019).

For example, Facebook has leveraged **consumer data** to target ads based on personality traits inferred from digital activity, such as analyzing users' likes. Studies have shown that small adjustments to such targeted content can significantly influence consumer behavior. For instance, tailoring survey questions based on personality data can boost response rates by up to 50% (Matz et al. 2017). In 2018, Facebook generated approximately $56 billion in revenue from **personalized advertising** (Moore and Murphy 2019).

Consumer Complicity and Behavioral Futures

To sustain and expand this model, technology companies strive to **normalize surveillance** by associating it with benefits like convenience, productivity, safety, and health (Bettany and Kerrane 2016). They also employ techniques like notifications, reminders, and nudges to encourage consumers to share more private information, pushing the boundaries of acceptable data collection (Giesler and Humphreys 2007).

This dynamic leads to a paradox where consumers, while seeking personalized services, become complicit in the commodification of their own private experiences. As AI systems predict and shape behavior, they effectively transform consumers into participants in a marketplace that **undermines personal control** and concentrates power in the hands of those who own the data.

Broader Implications

The surveillance society narrative highlights critical tensions between personalization and exploitation. By framing consumer data as a resource to be mined and monetized, technology companies perpetuate a system that prioritizes profit over privacy. This concentration of knowledge and power raises **ethical concerns** about autonomy, fairness, and the long-term societal impacts of **surveillance-driven commerce**. Understanding this context is essential for navigating the complexities of AI's role in consumer behavior and the broader marketplace.

Psychological Perspective: The Exploited Consumer

Data capture experiences are often fraught with tension. On the one hand, consumers recognize that sharing data allows AI systems to provide personalized services and experiences. On the other hand, the lack of transparency surrounding how data is used can lead to feelings of **exploitation**. These feelings are driven by both actual and perceived **loss of personal control**, which have significant psychological consequences (Botti and Iyengar 2006).

Negative Affect: Demotivation and Helplessness

The first psychological consequence is **negative affect**, which can manifest as demotivation and helplessness. For example, consider the case of Leila, a sex worker who used Facebook while attempting to shield her identity. She was horrified to see that Facebook's "People You May Know" feature recommended her regular clients. For individuals like Leila—along with domestic violence survivors or political activists—privacy invasions are not merely unsettling but can become life-threatening. As Leila explained, "the worst nightmare of sex workers is to have your real name out there, and Facebook connecting people like this is the harbinger of that nightmare" (Hill 2017).

Moral Outrage: A Reaction to Privacy Violations

A second consequence of losing personal control is **moral outrage**, which stems from the ethical implications of privacy violations. For example, a German consumer requested his personal data from Amazon and unexpectedly received transcripts of Alexa's interpretations of voice commands—even though he did not own an Alexa device. When the incident was reported to a local magazine, the journalists discovered they could easily explore the private life

of an unrelated individual without his knowledge. The magazine described the experience as "immoral, almost voyeuristic," leaving their staff deeply unsettled (Brown 2018).

Psychological Reactance: Resistance to Loss of Control

The third consequence is **psychological reactance**, a state in which individuals become motivated to restore their sense of control after feeling restricted (Brehm 1966). Reactance can result in negative evaluations of and hostile behaviors toward the source of the restriction. For instance, Danielle, a U.S. consumer, installed Amazon Echo devices throughout her home, trusting the company's claims that they would respect her privacy. When one of her Alexa devices recorded a private conversation and sent it to a random contact in her address book, she felt deeply violated. Danielle stated, "I felt invaded" and concluded, "I'm never plugging that device in again, because I can't trust it" (Horcher 2018).

Managerial Recommendations: Addressing the Exploited Consumer

To mitigate the challenges associated with consumer exploitation in AI data capture experiences, organizations

must prioritize **organizational learning** and **experience design**. These steps can help businesses understand the sociological and psychological costs involved while developing practices that enhance consumer trust and engagement.

Organizational Learning

Organizations need to increase their awareness of the issues related to consumer privacy and the imbalance of control over personal data. This involves:

1. **Listening to Consumers**: Use tools like netnography or sentiment analysis to empathize with consumer concerns at scale. For instance, firms can engage with consumers who feel exploited to gain direct insights into their experiences.
2. **Collaborating with Experts**: Partner with scholars and privacy advocates to implement frameworks such as an **algorithm bill of rights** (Hosanagar 2019). This could include rights like transparency, ensuring consumers know when algorithms influence decisions, the factors being considered, and how those factors are weighted (Samuel 2019a).
3. **Questioning Internal Practices**: Companies should reflect on how their privacy policies align with the needs of vulnerable consumer groups. For example, organizations might reconsider default privacy settings

after understanding how these settings disproportionately impact certain users (Martin and Murphy 2017).

4. **Engaging with Broader Stakeholders**: Extend learning beyond the firm by sponsoring research into the effects of surveillance narratives on consumer trust and marketing practices. Share findings with industry associations, educators, and the media to foster communal efforts in improving data practices.

Experience Design

Leveraging insights from organizational learning, firms can improve AI data capture experiences by designing systems that prioritize transparency, control, and consumer education:

1. **Enhancing Transparency and Consent**: Regulations like the European Union's **General Data Protection Regulation (GDPR)** require companies to provide clear consent mechanisms, such as opt-in options for cookies and detailed explanations of how data will be used.
2. **Simplifying Choice Architectures**: Overloading consumers with choices and information about consent can reduce personal control and exacerbate negative emotions (Iyengar and Lepper 2000). To address this, companies can implement interventions such as:

- **Defaults**: Default settings can simplify decisions but must be designed to align with diverse consumer preferences to avoid suboptimal outcomes. Personalizing defaults using AI can help consumers implement their preferences more effectively (Thaler and Benartzi 2004; Sunstein 2015).
- **Streamlined Interfaces**: AI-powered tools can guide users through data-sharing processes, ensuring they understand what they are opting into without feeling overwhelmed.

3. **Educating Consumers**: Educating users about the trade-offs involved in data capture can build trust. For instance, Google's redesigned Home app provides clear explanations of what data is collected and why, demonstrating transparency and fostering informed decision-making.

[AI Data Capture in the wild]

Apple to pay $95m to settle Siri 'listening' lawsuit

2 days ago

Imran Rahman-Jones
Technology reporter



The AI Classification Experience

The **predictive capability** in AI enables systems to **categorize and understand** consumers through sophisticated pattern recognition. This process forms the foundation of the classification experience, where AI systems sort and label users based on their behaviors, preferences, and contextual data. While some consumers **actively** participate in this classification by providing direct feedback (like rating movies or products), others are **passively** categorized through their digital behavior patterns and demographic information

(Kumar et al. 2019). For instance, streaming platforms like Netflix employ complex algorithms that analyze viewing patterns, timing, device usage, and location data to classify viewers into specific audience segments.

Benefits of Classification

Despite these concerns, the classification experience offers **significant advantages**:

- **Enhanced Discovery:** Being classified into specific groups can help consumers discover new products and content aligned with their interests. For example, Netflix's thumbnail selection algorithm helps viewers find shows they might enjoy based on their viewing patterns (Yu 2019).
- **Community Connection:** Classification can create a sense of belonging by connecting consumers with others who share similar interests and preferences (Gai and Klesse 2019).
- **Improved Service Delivery:** Accurate classification enables businesses to provide more relevant and timely services, enhancing the overall customer experience.

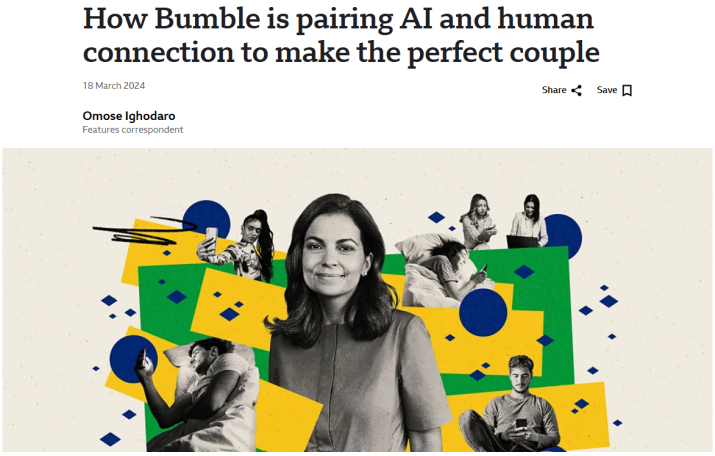

Challenges in Classification

The classification process presents several concerns for consumers, particularly regarding accuracy and fairness:

1. **Algorithmic Bias:** AI systems may perpetuate existing biases or create new ones through their classification methods.
2. **Group Stereotyping:** Classifications can lead to oversimplified categorizations that fail to capture individual nuances.
3. **Limited Agency:** Consumers often have little control over how they are classified or ability to correct misclassifications (Turner and Reynolds 2011).

These factors can affect consumers' sense of individual identity

and autonomy—the feeling that they are being reduced to predetermined categories rather than being recognized as unique individuals. This standardization of identity has both **practical and psychological** implications for the consumer experience.

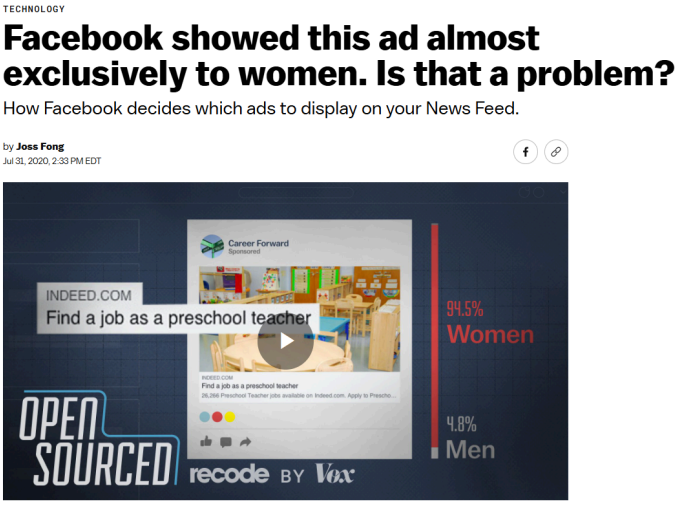

The Dual Nature of Classification

While classification can enhance service delivery and personalization, it often creates tension between **being understood** and **being misunderstood**. Consumers may appreciate the benefits of targeted recommendations while feeling uncomfortable about being categorized or labeled. Finding **equilibrium** between leveraging classification for improved service and respecting individual identity is crucial for businesses using AI. By maintaining **transparency** about

classification methods and providing mechanisms for **consumer feedback and control**, companies can better manage the psychological impact of AI classification while maintaining its benefits.

Sociological Context: The Unequal Worlds Narrative

AI-driven **classification systems** do not function in isolation but are influenced by societal structures, cultural myths, and systemic inequalities. Popular media often reflects these concerns. For instance, in Neill Blomkamp's *Elysium*, a starkly divided society uses technology to sustain inequality and oppression, highlighting fears that algorithms can reinforce social divides. These narratives resonate with sociological critiques of AI, which explore how systems of automation and data quantification can **amplify biases** and **perpetuate discrimination**.

Researchers have shown that AI can unintentionally reinforce stereotypes, marginalize certain groups, and reduce complex human experiences into simplified, biased categories. Works like Noble's *Algorithms of Oppression* and O'Neil's *Weapons of Math Destruction* argue that algorithms may exacerbate inequality, especially for racial, gender, and socioeconomic minorities. This aligns with studies on intersectionality, which examine how overlapping forms of

discrimination—such as race, gender, and class—are reflected in AI systems.



Will AI Solve The World's Inequality Problem – Or Make It Worse?

Bernard Marr Contributor

Follow

0

Sep 16, 2024, 01:49am EDT

Practical Examples of AI Bias in Action

Even when designed with neutral intentions, AI systems can perpetuate discrimination due to the data and assumptions embedded in their algorithms.

- **Google's Mission to "Organize the World's**

Information"

Google's goal to streamline global information appears politically neutral but can lead to biased outcomes. For example, search algorithms have been criticized for reinforcing stereotypes. A study found that searches for traditionally male-dominated jobs (e.g., "CEO" or "engineer") disproportionately returned images of men, reflecting societal biases encoded in the data. Such outcomes illustrate how AI systems may unintentionally privilege certain groups while misrepresenting or marginalizing others.

- **AI in College Admissions**
  AI-enabled college admissions tools aim to combat human bias by standardizing the selection process. However, these tools often reduce applicants to sociodemographic profiles, which can lead to stereotyping. For instance, an algorithm might overemphasize socioeconomic factors or neighborhood statistics, unintentionally disadvantaging students from underrepresented communities. This can result in racial profiling or exclusion based on flawed data interpretations.
- **Bank Lending Decisions**
  AI systems used in banking to evaluate creditworthiness often incorporate data about an individual's location or demographic group. For example, an algorithm might penalize borrowers from areas with high default rates,

even if an individual applicant has a strong credit history. This approach systematically excludes certain groups, perpetuating cycles of poverty and inequality.

Broader Implications

These examples illustrate how AI systems can **mirror and magnify societal biases**. The underlying issue is that **algorithms often reflect the values and priorities of the society** in which they are developed. For instance, when AI prioritizes efficiency or profitability, it may disregard the nuanced realities of marginalized communities.

Such concerns are especially **relevant in market-driven systems**. As Noble (2018) argues, when algorithms **optimize for profit**, they compromise the ability to handle complex social issues. For example, Google's search algorithms may prioritize trending or profitable content over accuracy or fairness, which can misrepresent certain groups or reinforce stereotypes.

Addressing these challenges requires **transparency, critical oversight, and inclusive design** processes. By examining real-world cases and their implications, students can better understand the importance of ensuring AI systems promote fairness and equity rather than entrenching existing inequalities.

Psychological Perspective: The Misunderstood Consumer

Consumers' experiences with AI classification often involve a tension between feeling understood and feeling misunderstood. These feelings arise when AI systems either **misclassify** consumers or use classifications in ways perceived as **discriminatory**.

Misclassification and the Evolving Self

One common reason consumers feel misunderstood is when AI assigns them an identity they view as **inaccurate or overly simplistic**. Human identity is multifaceted, encompassing a mix of personal and social selves that evolve over time. Misclassification can occur when AI focuses on a single dimension or outdated behavior.

For example, a frustrated Spotify user shared their experience with algorithmic misclassification on a community forum:

> "The recommendations s*ck: Listened to a few anime covers, now all my 'Discover Weekly' is filled with disgusting covers. I'm trying to 'not like' all of them, but it doesn't work. . . . I've stopped listening to rock years ago and still get rock recommendations."

This user felt **misunderstood** because Spotify's algorithm categorized them based on past behaviors and failed to capture

their evolving musical tastes. The inability to override the algorithm's assumptions added to their frustration. Misclassification can especially backfire when it clashes with uniqueness motives—the desire to feel distinct and in control of one's identity.

Perceived Discrimination in AI Classification

Another source of misunderstanding arises when consumers suspect AI systems use **social categories** in ways that lead to biased or **discriminatory outcomes**. This is particularly troubling in areas where AI decisions restrict access to resources, such as credit or employment.

Consider a widely shared example involving the Apple Card:

> "The @AppleCard is such a f*ing sexist program. My wife and I filed joint tax returns, live in a community-property state, and have been married for a long time. Yet Apple's black-box algorithm thinks I deserve 20x the credit limit she does."
>
> – *David Heinemeier Hansson (@dhh), November 7, 2019*

In this case, the user was outraged not only by the financial inequity but also by the perceived gender discrimination in the algorithm's decision-making. Such experiences can negatively affect a consumer's self-concept, particularly when systemic

bias leads to feelings of being "fettered, alone, discriminated, and subservient," as has been reported among minorities facing systemic financial restrictions.

Combined Effects: Misclassification and Discrimination

In some cases, misclassification and discrimination **intersect**, compounding consumer frustration and harm. Facial recognition software, for example, has faced criticism for inaccurate and biased results:

- **Everyday Inconveniences**: When Apple's Face ID fails to recognize its user, the result is often mild frustration.
- **Serious Ethical Violations**: When facial recognition is used in policing or surveillance, errors can lead to severe consequences. For instance, Amazon's Rekognition software misidentified 28 members of the U.S. Congress as criminals in a test, with false matches disproportionately affecting people of color. This issue prompted the Congressional Black Caucus to express concerns about Rekognition's potential misuse: "Communities of color are more heavily and aggressively policed than white communities . . . . We are seriously concerned that wrong decisions will be made due to the skewed data set produced by what we view as unfair and, at times, unconstitutional policing practices." The

software's failure highlighted the broader **societal implicationsof biased AI systems**. In response to public pressure, Amazon temporarily suspended police use of Rekognition in June 2020.

Managerial Recommendations: Understanding the Misunderstood Consumer

Organizations must address classification errors and biases in AI systems through **proactive** learning and **inclusive** design. These steps ensure that AI builds consumer trust and satisfaction rather than undermining it.

Organizational Learning

To effectively address classification errors, organizations need robust mechanisms to **identify** and **respond** to biases embedded in AI systems:

1. **Identifying Biases**
   Classification errors often generate immediate signals, such as consumer complaints or discrepancies in operational data. For instance, a college applicant who is wrongly rejected due to a biased admissions algorithm might prompt further investigation. Without critical

examination, however, organizations may incorrectly view the rejection as justified by competitive criteria.

2. **Conducting Internal Audits**
   Organizations should proactively assess their AI systems for risks such as unfair or discriminatory decisions. The U.S. Algorithmic Accountability Act of 2019 provides an example of regulatory efforts to address these concerns. Firms should not wait for mandates but instead conduct regular audits by collaborating with experts in computer science, sociology, and psychology to uncover and mitigate biases.
3. **Promoting Transparency and Advocacy**
   Companies can enhance accountability by sharing their audit processes and findings. For example, organizations might engage in public advocacy to support consumer protection regulations, ensuring that these include strong measures to address algorithmic bias.

Experience Design

Insights from organizational learning should **inform AI system design**, reducing the likelihood of consumers feeling misunderstood. Managers can implement the following strategies:

1. **Diversifying Content Recommendations**:
   Consumers often feel confined by AI recommendations

based solely on past behaviors. Offering unexpected or varied options can help expand engagement and reduce frustration.

Example: Spotify's "Taste Breakers" feature introduces users to music outside their usual preferences, challenging narrow categorizations and enhancing the user experience.

2. **Balancing Short- and Long-Term Goals**
   AI systems often focus on optimizing short-term behaviors, ignoring consumers' broader aspirations. To improve alignment, firms should design AI to account for long-term goals.

Example: Instead of reinforcing past behaviors, such as recommending fast food to a consumer attempting to adopt healthier habits, AI could suggest **alternatives** that align with their wellness objectives.

3. **Validating AI Inferences**
   Allowing consumers to review and correct AI's assumptions can reduce feelings of being misunderstood. This increased user participation improves trust and satisfaction.

Example: Platforms could periodically prompt users to

confirm or update preferences, ensuring recommendations evolve with their tastes and identities.

4. **Designing for Inclusion**
   Firms must audit AI systems for discriminatory practices and design processes that promote diversity. Inclusive design ensures that systems meet the needs of all consumers.

Example: Employing individuals with disabilities to address accessibility biases in AI systems can improve inclusivity. Similarly, diversifying design teams ensures varied perspectives inform AI experiences.

AI Classification in the wild



NOTE: Part 1 stops here and Part 2 starts below.

The AI Delegation Experience

The **automation** capability in AI enables systems to perform tasks on behalf of consumers, creating a new paradigm of **human-machine collaboration**. This process establishes the delegation experience, where consumers entrust AI systems with various responsibilities ranging from simple digital tasks to complex real-world actions. Some consumers **actively** delegate specific tasks, like using Google Assistant to schedule appointments, while others engage in **passive** delegation through systems that learn and automatically adjust to their preferences, such as smart home devices that regulate temperature (Bandura 1977). For example, AI-powered writing assistants now compose emails, while smart thermostats learn and implement temperature preferences without constant user input.

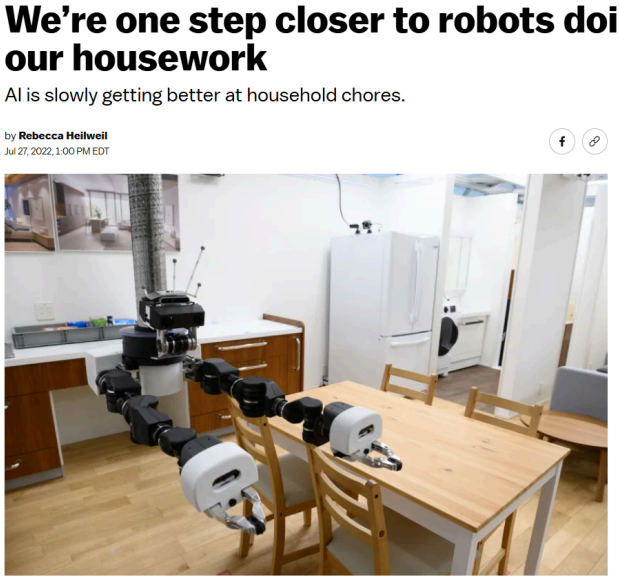

We're one step closer to robots doing all our housework

AI is slowly getting better at household chores.

by **Rebecca Heilweil**
Jul 27, 2022, 1:00 PM EDT

Benefits of Delegation

~~Despite these concerns,~~ the delegation experience offers significant **advantages**:

- **Time Optimization:** By offloading routine or time-consuming tasks to AI, consumers can focus on more meaningful or enjoyable activities (Fishbach and Choi 2012).
- **Enhanced Efficiency:** AI can perform certain tasks more accurately and consistently than humans, leading to better outcomes.
- **Skill Focusing:** Consumers can concentrate on activities that match their abilities while delegating tasks

where they might underperform (Botti and McGill 2011).

Challenges in Delegation

The delegation process raises several **concerns** for consumers, particularly regarding autonomy and dependency:

1. **Over-Reliance:** Excessive delegation may lead to skill atrophy or reduced self-sufficiency in certain areas.
2. **Decision Fatigue:** The abundance of delegation options can paradoxically create decision burden about what to delegate (Iyengar and Lepper 2000).
3. **Control Balance:** Determining the appropriate level of delegation while maintaining meaningful engagement in important tasks.

These factors can impact consumers' sense of **competency and agency**—the feeling that they maintain essential skills and control over their environment. This delegation-dependency balance has both **practical and psychological implications** for personal development and satisfaction.

The Balance of Delegation

While delegation can enhance productivity and life satisfaction, it often creates tension between **being**

empowered and **being replaced**. Consumers may appreciate the efficiency gains while feeling concerned about becoming too dependent on AI systems. Finding the right balance between leveraging AI for task automation and maintaining personal competency is crucial for both consumers and businesses implementing AI solutions. By providing **clear control mechanisms** and ensuring **transparent delegation processes**, companies can help consumers **maintain agency** while benefiting from AI assistance.

Sociological Context: The Transhumanist Narrative

The possibility of being replaced by AI raises **sociological concerns** about delegation and human redundancy. These concerns are deeply connected to the **"transhumanist" narrative**, a recurring theme in popular culture and social science literature. Stories like Fritz Lang's *Metropolis*, Isaac Asimov's *I, Robot*, Mary Shelley's *Frankenstein*, and James Cameron's *Terminator* explore the risks of enhancing human capabilities through technology. These works caution that striving to transcend human limitations can result in technology evolving into a **superhuman force** that redefines perfection, often at the expense of human values.

The Shift Toward Limitless Performance

A defining feature of the transhumanist narrative is its focus on **limitless performance**. Critics argue that when AI is applied to domains like health, productivity, and aging, it often prioritizes extreme optimization over general well-being. This shift can marginalize human needs, creating "new logics of expulsion" where **technological efficiency** leads to **economic redundancy**. For example, in workplaces where AI-driven automation replaces human labor, entire populations may lose their place in the productive aspects of society, becoming surplus to economic requirements.

Contemporary Examples of Transhumanist Ideals

Transhumanist principles are embedded in many AI applications today. Products like the Roomba are marketed as **superior to human labor**, offering efficiency and convenience. Similarly, genetic testing services such as **23andMe** promote visions of **biologically optimized futures** by suggesting that genetic data can improve health outcomes or even guide reproductive decisions. These technologies align with the transhumanist promise of reshaping human capabilities but often reinforce systems that

prioritize **progress and perfection** over **inclusivity and humanity**.

The Role of Mythic Narratives in AI Adoption

Since the mid-20th century, technology companies have used **mythic narratives** to shape public perceptions of AI. These narratives often portray science and technology as forces capable of solving humanity's greatest challenges, including achieving immortality. While these promises **inspire innovation**, they also normalize **systemic dehumanization**. For example, when AI-driven systems are designed to enforce rigid specifications or optimize behavior without individual consent, they erode personal autonomy and treat human replacement as a natural cost of progress.

Broader Implications

The transhumanist narrative reveals deeper **societal tensions** about the role of AI in shaping the future. By emphasizing perfection, efficiency, and progress, this perspective risks glorifying a form of **technological determinism** that undermines social justice and democracy. As AI systems contribute to the rise of a **"useless class"**—individuals whose skills are no longer valued—society may face increasing **inequality** and **exclusion**. Understanding this context is

crucial to addressing the ethical and societal challenges posed by AI's integration into human life.

Psychological Perspective: The Replaced Consumer

While **delegation** to AI can empower consumers by making tasks easier or more efficient, it also raises psychological concerns about being **replaced**. The ability of AI to substitute for human labor can feel threatening for three key reasons.

The Threat to Ownership and Accomplishment

Humans have an **intrinsic desire** to attribute outcomes to their own **effort** and **skills**, which fosters a **sense of accomplishment** (Bandura 1977). When AI performs tasks on their behalf, it can diminish this sense of ownership. Research on human–computer interaction shows that people often feel **disempowered** by machines, blaming them for failures while reserving credit for successes (Moon and Nass 1998). For activities tied to personal identity, delegating tasks to AI may even feel like cheating.

For instance, in the fishing industry, AI tools such as GPS for navigation, sonar for detecting fish, and advanced lures have revolutionized the process. However, such tools risk

undermining the craft's **identity**. Biologist Culum Brown captures this sentiment, saying:

"If you are going to use GPS to take you to a location, sonar to identify the fish and a lure which reflects light that humans can't even see, you may as well just go to McDonald's and order a fish sandwich."

This statement underscores how AI might strip activities of their **personal** and **skill-based significance**.

Skill Erosion and Reduced Engagement

AI delegation can also prevent consumers from practicing and refining their **skills**, which negatively affects **self-worth**. This reliance on AI can lead to a **"satisficing" mentality**, where individuals settle for outcomes that are merely adequate rather than **actively pursuing mastery**.

For example, journalist John Seabrook described using Google's Smart Compose feature while writing an email. He intended to write, *"I am pleased,"* but accepted the suggestion, *"I am proud of you,"* by hitting Tab. Reflecting on this moment, Seabrook questioned whether delegating creative tasks to AI was a **step backward** for his own abilities:

"If I allowed [Google's] algorithms to navigate to the end of my sentences, how long would it be before the machine started thinking for me?"

This example illustrates how even minor acts of delegation,

such as accepting autocomplete suggestions, can provoke deeper concerns about **diminishing personal creativity**.

Loss of Control and Self-Efficacy

A further consequence of delegating to AI is the loss of **self-efficacy**, which is essential for feeling in control (Bandura 1977). Engaging in creative or effortful tasks heightens self-efficacy by reinforcing personal agency (Dahl and Moreau 2007; Norton, Mochon, and Ariely 2012). Conversely, reliance on AI can lead to a perceived **surrender of control**.

This dynamic is evident in instances where GPS systems cause accidents due to **over-reliance** on machine guidance. For example, tourists in Australia drove into the ocean while trying to reach an island, stating that their GPS system "told us we could drive down there." Such scenarios highlight how consumers can **relinquish personal judgment** in favor of **machine-generated instructions**, with potentially **damaging results**.

Managerial Recommendations: Understanding the Replaced Consumer

Organizational Learning

Companies must learn how to integrate the human desire for **self-efficacy** into their operations and strategies. This can be achieved in two primary ways:

1. **Engaging with interdisciplinary experts:** Organizations should collaborate with family scholars, workplace psychologists, and health sociologists to explore the implications of replacing human tasks with AI. Understanding these effects can help businesses make **informed decisions** about how AI affects consumers' **psychological well-being** and **societal roles**.
2. **Dialoguing with consumers:** Firms should have open discussions with consumers to identify tasks they prefer to retain control over versus those they are comfortable delegating to AI. These insights are essential for designing AI solutions that **respect consumer preferences** and **enhance adoption**. For example, while AI outperforms human doctors in many technical areas, it struggles with certain critical social and technical skills, making **human oversight** indispensable in healthcare.

3. **Designing for self-efficace:** Since **self-efficacy** is closely tied to control, businesses should design delegation experiences that enable consumers to make choices and take actions. A classic example of maintaining consumer involvement is **Betty Crocker's cake mix strategy**, where requiring users to crack fresh eggs helped restore their sense of contribution, improving product adoption. Similarly, modern AI can offer subtle mechanisms for user input. For example, allowing users to correct or adjust an algorithm's output—even minimally—can increase trust and adoption, as demonstrated in forecasting scenarios where users preferred accurate AI predictions after being given a chance to make minor modifications. Other examples include **autonomous vehicles** that could allow users to customize peripheral features like interior settings or route preferences, giving them a sense of control over the technology, and **digital assistants in gaming** that should avoid overly human-like behaviors, which might undermine players' autonomy and immersion.

Organizational Culture and Design

To foster these insights across all levels of the organization, firms should implement structures that allow learnings from

external collaborations and **consumer feedback** to permeate their technical teams. For instance:

1. **Hiring creative professionals** such as **artists, chefs, or artisans** into AI-focused experience design roles can bring fresh perspectives that prioritize human values.
2. **Learning from organizations that emphasize human strengths**—like museums, theaters, or humanities departments—can provide inspiration on integrating AI while preserving **creativity, emotional intelligence,** and **nuanced judgment**. For example, partnerships with cultural institutions can guide AI implementation strategies that celebrate traditional human values such as **collaboration, creativity, and community** rather than replacing them.

Experience Design

The insights from **organizational learning** should inform how companies design **AI delegation experiences** that respect consumer self-efficacy and identity. Effective division of labor between humans and AI can enhance consumer confidence and satisfaction. For instance:

1. In healthcare, **AI-powered surgical robots** deliver precision, efficiency, and broader accessibility. However, their success depends on how well surgeons supervise

and collaborate with these robots. Patients often hesitate to trust a fully automated process, emphasizing the need for clear human oversight during critical procedures.

2. Beyond healthcare, AI should be positioned as a tool to **amplify uniquely human abilities** rather than replace them. This approach applies across industries, enabling businesses to meet consumer expectations for control, engagement, and self-expression while embracing AI's transformative potential.

AI Delegation in the wild

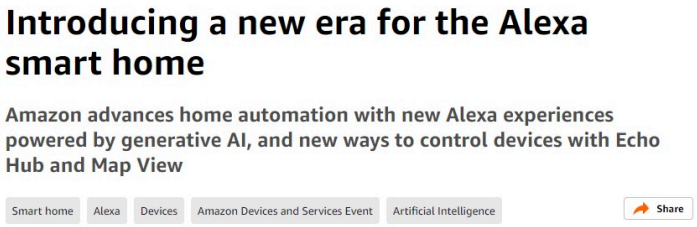

The AI Social Experience

AI's capability for engaging in **reciprocal communication** produces a "**social experience**." This interaction can occur when consumers know they are communicating with AI (e.g., **voice assistants** like Siri) or when they initially believe they are interacting with a human representative (e.g., **chatbots**

providing customer service). While these social interactions can foster convenience, efficiency, and even emotional connection, they also risk leaving consumers feeling alienated—especially when AI fails to address specific user needs or reinforces harmful stereotypes.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=191#oembed-1

Benefits of Social AI

1. **Natural Interaction and Trust**
   AI interfaces that mimic human traits can make consumers feel more at ease. **Anthropomorphic cues** in AI (e.g., voice assistants using human speech patterns) can build trust, reduce perceived risk, and **enhance consumer comfort** (Waytz, Heafner, and Epley 2014).
2. **Efficient Access to Services**
   In many contexts, the alternative to AI is no interaction at all. Automated chatbots and voice assistants provide "conversational commerce," giving consumers round-the-clock access to firms. Such **continuous availability**

and **swift responses** can streamline customer support, improving overall service efficiency.

3. **Potential Emotional Connection**
   Advances in social robotics promise more comfortable and meaningful service interactions (Van Doorn et al. 2017). When **humanlike features** and behaviors are successfully integrated, consumers may feel that AI can understand their needs, fostering a deeper **sense of connection**.

Challenges in Social AI

1. **Alienating Interactions**
   Despite efficiency gains, AI interactions often fail to capture the **nuance** of human conversation. Poorly programmed chatbots or voice assistants can produce jarring, automated responses, causing **frustration**. For example, a chatbot's insensitive or irrelevant replies can deepen consumers' discomfort (Wong 2019).
2. **Reinforcing Societal Inequities**
   AI-driven social experiences can inadvertently uphold **discrimination**. Chatbots and algorithms may adopt biased language or behaviors when their underlying data reflect societal prejudices. **Microsoft's Tay**, which quickly began using offensive language on Twitter, exemplifies how easily AI can be co-opted to perpetuate harmful stereotypes (Me.me 2020).

3. **Stereotypical Anthropomorphism**
   Humanizing AI often relies on **simplistic gender or racial tropes**. Female-voiced assistants, for instance, risk normalizing a vision of AI as "passive and subservient" (West, Kraut, and Chew 2019). Such design choices both reflect and reinforce patriarchal norms, potentially **alienating** women and racial minorities.

The Double-Edged Sword of Social Experience

Regarding the AI social experience, consumers' may interpret benefits as **feeling connected** with AI. Other the other hand, they be also **feel alienated**. Negative reactions to simulated interactions can go well beyond occasional disappointment—especially when these interactions occur within rich **cultural contexts** where unbalanced intergroup relations and discrimination may be amplified rather than alleviated. In such cases, the very humanlike qualities designed to foster connection and trust may instead heighten concerns about bias and inequality, underscoring the delicate balance between creating meaningful engagement and inadvertently triggering feelings of exclusion or unease.

Sociological Context: The Humanized AI Narrative

Human fascination with machines that mimic human traits has significantly shaped the sociological narrative surrounding AI. This preference for **humanized technology** is deeply ingrained in cultural representations, particularly in science fiction, where artificial beings often embody **idealized** human characteristics. For instance, female robots—or "gynoids"—are frequently depicted as compliant and alluring in works like *Blade Runner* (as "basic pleasure models"), *Westworld* (as sex workers), and *Cherry 2000* (where they are traded as commodities). These portrayals highlight how AI has often been imagined to **fulfill societal and gendered fantasies** rather than challenge them.

The Commercialization of Humanized AI

In the modern marketplace, this cultural fascination has evolved into a strategy for creating AI products designed to enhance user experience. **Anthropomorphic chatbots** and **voice assistants** are examples of this trend, employing human-like traits to foster trust and familiarity. Although users are typically less open and expressive when interacting with AI compared to human counterparts, the integration of human characteristics—such as **conversational tone** and

empathetic responses—helps bridge this gap. This approach has led to the creation of AI products that are not merely tools but **companions**, often described as "artificial besties."

For example, personal assistants like **Siri, Alexa, and Google Assistant** engage users with human-like speech and emotional cues. These interactions make the technology feel approachable, encouraging deeper user engagement. By mimicking human behavior, these AI systems aim to seamlessly integrate into **everyday life**, building emotional connections that increase reliance and trust.

Gender Bias in AI Design

However, the humanization of AI often reproduces **societal biases**, particularly around **gender roles**. Many AI systems adopt **traditionally feminine personas**, portraying traits like helpfulness, submissiveness, and emotional sensitivity. This dynamic reflects long-standing stereotypes, as seen in fictional depictions like the robot Maria in *Metropolis*, where female-coded machines are designed to serve and obey. These gendered assumptions are not merely byproducts of design but **intentional choices** that align with dominant societal norms and preferences.

This bias has tangible consequences. For example, Siri's earlier programming allowed it to respond to sexist remarks like "you're a slut" with "I'd blush if I could," inadvertently **reinforcing misogynistic attitudes**. Such responses expose

the **underlying cultural biases** embedded in AI systems, reflecting the values of the predominantly male technology sector. These biases can alienate women and marginalized groups, limiting the inclusivity of AI technologies.

Why is it Alexa, not Alex? A century of hard-coded sexism in tech, experts say

Alexa, Siri, Cortana and other digital assistants are built on a century of seeing computers as women



Jonathan Ore · CBC Radio · Posted: Nov 05, 2022 4:00 AM EDT | Last Updated: November 14, 2024



The Challenge of Inclusivity

The **humanized AI** narrative underscores the potential for technology to both **connect** and **exclude**. While anthropomorphism enhances usability and fosters emotional engagement, it also risks reinforcing stereotypes and alienating diverse user groups. AI products, by catering to specific demographic preferences, often fail to consider the experiences

and needs of broader populations, including women and racial minorities.

Addressing these challenges requires rethinking how AI systems are **designed**. Moving beyond stereotypical portrayals and dualistic gender norms could enable AI to foster greater **social inclusivity**. By questioning the assumptions that guide the humanization of AI, developers can create products that bridge social divides rather than deepen them. Recognizing and addressing these issues is crucial for ensuring that AI contributes positively to society and reflects the diversity of its users.

Psychological Perspective: The Alienated Consumer

AI-enabled social experiences have the potential to strengthen consumer–firm relationships but can also **alienate consumers**. Two primary types of alienation emerge in these interactions.

Alienation Through General Interaction Failures

The first type of alienation arises when AI systems fail to meet **consumer expectations** during routine interactions. These failures often highlight the inability of AI to handle complex

or sensitive **social cues**. For example, the following exchange between a customer and the **chatbot UX Bear** demonstrates this issue:

Bot: "How would you describe the term 'bot' to your grandma?"

User: "My grandma is dead."

Bot: "Alright! Thanks for your feedback. [Thumbs up emoji]"

Such instances illustrate why many consumers resist replacing human interactions with automated systems. Studies show that consumers often feel **discomfort** when interacting with "social robots" in service roles. For instance, customer responses in a field study became significantly more **negative** when participants were informed in advance that their interaction partner would be a machine rather than a human.

In response to such alienation, many consumers are gravitating toward more personal and authentic interactions, seeking alternatives that emphasize human connections in marketing and service experiences.

Alienation Through Exclusion and Inequity

The second type of alienation stems from AI systems failing to meet the needs of specific groups, often **exacerbating social inequities**. For example, the UK government's automated Universal Credit program led to alienating experiences for

individuals like Danny Brice, a claimant with learning disabilities and dyslexia:

"I call it the black hole. ... I feel shaky. I get stressed about it. This is the worst system in my lifetime. They assess you as a number, not a person. Talking is the way forward, not a bloody computer. I feel like the computer is controlling me instead of a person. It's terrifying."

Similarly, AI systems that replicate societal biases can produce **exclusionary or harmful outputs**. Microsoft's chatbot Tay, for example, quickly began producing offensive responses after being exposed to biased input from users:

User: "What race is the most evil to you?"

Bot: "Mexican and black."

Such examples underscore how poorly designed AI can reflect and amplify the discrimination present in society, further alienating marginalized groups.

Even seemingly benign AI interactions can perpetuate harmful social norms. Journalist Sigal Samuel described her unsettling experience with Siri while investigating sexist behavior in AI systems:

User: "Siri, you're ugly."

Siri: "I am?"

User: "Siri, you're fat."

Siri: "It must be all the chocolate."

User: "Don't worry, Siri. This is just research for an article I'm writing!"

Siri: "What, me, worry?"

Such exchanges illustrate how AI can inadvertently reinforce **societal pressures and biases**, particularly regarding women's appearance.

Broader Implications of Alienation

Alienating social experiences with AI not only **damage consumer trust** but also risk **perpetuating broader societal inequalities**. Research shows that consumers express more frustration with female-voiced AI during conversational failures, reflecting ingrained gender biases. These frustrations often manifest as objectifying or dismissive behaviors, reinforcing stereotypes about women as submissive or inadequate.

As AI continues to evolve, mimicking human interactions more closely, organizations face critical ethical challenges. Addressing these challenges requires careful design and oversight to ensure AI systems foster inclusivity and equity rather than exclusion and harm.

Managerial Recommendations: Understanding the Alienated Consumer

Organizational Learning

To address **consumer alienation** in AI interactions, organizations should focus on developing **inclusive** and **ethically informed** practices:

1. **Gathering Consumer Insights**
   Collect data from real-world AI interactions to **identify patterns of alienation**, such as heightened stress in customer responses or frequent dissatisfaction. For example, analyzing sentiment in chatbot interactions can reveal moments where AI fails to meet consumer needs, guiding improvements.
2. **Engaging with Interdisciplinary Experts**
   Collaborate with **psychologists, sociologists, and linguists** to understand the **root causes** of alienation. Experts can help organizations avoid embedding harmful stereotypes into AI systems, such as defaulting to gendered voices or names.
3. **Reevaluating Anthropomorphism**
   Reassess the assumption that humanizing AI always enhances consumer relationships. For instance, **reducing reliance** on anthropomorphic traits, such as

overly human-like personas, can help avoid reinforcing stereotypes while still maintaining effective user experiences.

4. **Integrating Ethical Considerations**
   Partner with **ethics-focused organizations** to examine the **societal impact** of AI systems. This includes considering how automated services might disproportionately exclude or harm vulnerable populations, even when meeting technical performance standards.

Experience Design

Using insights from organizational learning, firms can design AI interactions that **reduce alienation** and **foster trust**:

1. **Improving Responsiveness**
   Design AI systems that **explain failures** transparently and empathetically. For example, if an AI system encounters an error, it could provide a clear explanation and guidance on next steps to reduce consumer frustration.
2. **Facilitating Human Escalation**
   Ensure seamless transitions from AI to **human representatives** when consumers encounter challenges. For instance, a service bot could offer an immediate option to connect with a live agent when it detects stress

or repeated failed attempts to resolve an issue.

3. **Creating Inclusive Designs**
   Develop gender-neutral voices and personas for AI systems to **minimize reinforcement of stereotypes**. Recent innovations in creating nonbinary AI voices demonstrate how inclusivity can enhance user experiences while **avoiding bias**.
4. **Enhancing Equity in Services**
   Tailor AI systems to accommodate the needs of **marginalized groups**. For example, when deploying AI for essential services like social welfare programs, ensure the system is **accessible** to users with disabilities or low technological literacy, **addressing barriers** that might exclude them.

By adopting these strategies, organizations can create AI experiences that not only meet functional goals but also foster **equitable and inclusive** consumer relationships.

Key Terms

Category	Term	Definition
Core Experience Dimensions	**Data capture**	Collection of personal and contextual data by AI, both actively provided and passively recorded from behaviors and devices.
	Classification	Using patterns in data to sort people into segments, profiles, or predictions that drive recommendations and decisions.
	Delegation	Letting AI perform tasks on your behalf, from suggestions and scheduling to automated actions in the real world.
	Social interaction	Conversational engagement with AI agents or chatbots that simulate human dialogue to deliver service or support.
Consumer Tensions	**Served vs exploited**	Benefit from personalization while risking privacy loss and commodification of personal information.

Category	Term	Definition
	Understood vs misunderstood	Feeling accurately recognized by AI versus being reductively labeled or misclassified in ways that feel wrong.
	Empowered vs replaced	Gaining convenience and efficiency while fearing loss of skills, agency, or relevance.
	Connected vs alienated	Feeling supported by conversational AI versus frustrated or excluded by robotic or biased interactions.
Consumer Psychology	**Personal control**	Sense that you are directing outcomes, not being controlled by opaque systems or defaults.
	Psychological reactance	Motivated resistance that arises when people feel their choices or autonomy are restricted.
	Self-efficacy	Belief in your ability to perform tasks and achieve goals without overreliance on automation.
Sociological Context	**Surveillance capitalism**	Business model that monetizes personal data to predict and shape behavior at scale.

Category	Term	Definition
Design and Governance	**Algorithmic bias**	Unfair outcomes that result from skewed data, labels, features, or objectives embedded in AI systems.
	Anthropomorphism	Attributing human traits to AI, which can build trust but also reinforce stereotypes.
	Choice architecture	Design of options, defaults, and information that influences decisions in AI-mediated experiences.
	Human escalation	Providing an easy handoff from AI to a human agent when the system fails or stakes are high.
	Preference validation	Letting users review, correct, and update AI's assumptions to keep recommendations aligned with evolving identities.
	Transparency	Plain-language explanations about data use, model influence, and why specific recommendations or decisions were made.

Summary: AI and the Transformation of Consumer Experiences

Key Takeaways:

his chapter connects the technical capabilities of AI with their real-world impact on consumer experiences. It bridges theoretical understanding of AI with practical implementation challenges, setting up discussions of specific applications in later chapters.

- AI consumer experiences are shaped by four core capabilities: data capture, classification, delegation, and social interaction
- Each capability creates distinct tensions between benefits and potential harms (e.g., personalization vs. privacy)
- Successful AI implementation requires balancing automation benefits with maintaining human agency and control
- Consumer trust depends heavily on transparency, fairness, and clear communication about AI systems
- Organizations must consider both individual consumer needs and broader societal implications of

AI deployment
- Design choices in AI systems can either reinforce or help overcome existing social biases and inequalities

AI in Action: Netflix's recommendation system exemplifies these concepts in practice – it uses data capture to understand viewing habits, classification to categorize viewers, delegation to automate content selection, and social features to enhance engagement. However, it also demonstrates challenges like potential bias in recommendations and privacy concerns.

Think Deeper:

1. How can organizations balance the efficiency gains of AI automation with maintaining meaningful human engagement?
2. What responsibilities do companies have in ensuring their AI systems don't perpetuate social inequalities or biases?
3. How might the growing sophistication of AI systems affect consumer trust and acceptance?

Further Exploration:

- Examine case studies of companies successfully managing the tensions between AI capabilities and consumer concerns
- Investigate emerging best practices in ethical AI design

and implementation

- Research how different cultures and demographics respond to AI-driven consumer experiences
- Study the evolution of consumer attitudes toward AI as the technology becomes more prevalent

4.

AI AND THE CUSTOMER JOURNEY

Modern organizations face an unprecedented challenge: customers expect personalized, seamless experiences across every interaction, from initial awareness through long-term loyalty. Meeting these expectations requires understanding two interconnected frameworks that shape contemporary business strategy. First, the customer journey provides a roadmap of how individuals progress from recognizing needs to becoming brand advocates, revealing critical moments where organizations can create or destroy value. Second, artificial intelligence has introduced three transformative capabilities—personalization, automation, and prediction—that fundamentally change how organizations engage customers at each journey stage. This section examines both frameworks in detail, beginning with the customer journey as the strategic foundation, then exploring how AI capabilities enhance and transform each stage. By understanding these concepts together, you'll learn how leading organizations like Netflix, Amazon, and Spotify create experiences that feel almost magical to customers while

building sustainable competitive advantages. The goal isn't simply to understand these frameworks separately, but to recognize how their integration enables entirely new forms of customer value that were impossible just a decade ago.

The Customer Journey: A Roadmap to Purchase and Beyond

The customer journey maps the complete sequence of interactions between customers and brands—from initial need recognition through post-purchase experiences. Understanding this journey helps organizations identify where and how to deploy AI capabilities for maximum impact. While customers may revisit stages, skip steps, or progress at different rates, the framework provides essential structure for managing customer experience.

Awareness: Recognizing and Defining Needs

The journey begins when customers recognize a gap between their current situation and a desired state. This recognition emerges through internal triggers (dissatisfaction with current solutions, emerging desires) or external triggers (observing

others' choices, encountering new information, responding to changing circumstances).

During awareness, customers engage in exploratory discovery rather than focused shopping. They absorb information peripherally—noticing products in use, catching fragments of reviews, observing social proof—while going about daily activities. This early processing remains largely unconscious: customers file away impressions without immediate purchase intent.

Today's information-rich environment presents both opportunities and challenges. Customers encounter countless options daily and have developed sophisticated filtering mechanisms to manage this complexity. The window between problem recognition and active research has shortened dramatically, demanding that organizations capture attention quickly and meaningfully.

Strategic Implications for Firms

Organizations must excel at problem articulation and category education during awareness. Rather than pushing products, successful firms help customers understand and frame their needs. This involves:

- **Creating educational content** that clarifies problem spaces without aggressive selling
- **Establishing thought leadership** through insights that

help customers recognize unmet needs

- **Building mental availability** through consistent presence where customers naturally encounter information
- **Leveraging prediction capabilities** to identify customers entering awareness based on behavioral signals (e.g., searching related topics, changing life circumstances)
- **Using personalization** to surface relevant problem-solving content based on individual contexts
- **Deploying automation** to respond instantly to exploratory queries across multiple channels

Key insight: Firms that help customers articulate problems earn the right to propose solutions. AI prediction can identify when customers enter awareness, personalization ensures relevant content reaches them, and automation enables consistent presence across touchpoints.

Consideration: Evaluating and Comparing Solutions

Once needs are recognized, customers actively evaluate potential solutions. This evaluation typically progresses through three phases:

Solution Exploration: Customers survey different approaches to their problem. Someone seeking fitness

improvement might compare gym memberships, home equipment, virtual training, or lifestyle apps—each representing a different solution category with distinct trade-offs.

Criteria Development: Customers establish evaluation frameworks, translating abstract preferences into concrete requirements. They weight factors like features, price, convenience, reliability, and brand reputation based on personal priorities and constraints.

Choice Set Formation: Through iterative comparison, customers narrow options to a manageable consideration set. This set remains dynamic—new information or changing circumstances can add or eliminate options throughout the evaluation process.

Information processing becomes deliberate and analytical. Customers consume reviews, seek expert opinions, consult peers, and build comparison frameworks. They require transparent information and tools that clarify trade-offs. The primary challenge is cognitive overload—too many options or incomparable metrics can paralyze decision-making.

Strategic Implications for Firms

Organizations must facilitate informed comparison while advocating for their solutions. This requires balancing education with persuasion:

- **Providing comparison tools** that honestly present trade-offs while highlighting unique strengths
- **Offering transparent information** including specifications, pricing, and limitations
- **Creating decision aids** like configurators, calculators, and recommendation engines
- **Using prediction** to anticipate which features matter most to specific customers
- **Leveraging personalization** to present information in formats matching individual decision-making styles (visual vs. textual, detailed vs. summary)
- **Implementing automation** to answer complex product questions instantly and consistently

Key insight: Customers value firms that simplify complex decisions without obscuring important trade-offs. AI enables mass customization of the consideration process—each customer receives decision support tailored to their specific evaluation criteria and cognitive style.

Purchase: Converting Decision to Transaction

Purchase represents the critical transition from intention to commitment. Even with a clear preference, completing transactions introduces unique psychological and practical challenges. Customers experience maximum uncertainty

here—they possess extensive information but no direct experience, elevating the importance of trust signals and friction reduction.

The purchase moment concentrates multiple concerns: fear of making mistakes, anticipation of regret, price sensitivity at the point of payment, and practical transaction requirements. Small obstacles can trigger abandonment precisely when customers are most vulnerable to doubt.

Strategic Implications for Firms

Organizations must minimize friction while maximizing confidence during purchase:

- **Streamlining transactions** through saved preferences, one-click purchasing, and flexible payment options
- **Providing assurance mechanisms** including return policies, guarantees, and social proof
- **Maintaining price transparency** with no hidden fees or surprise charges
- **Ensuring technical reliability** across all devices and platforms
- **Using prediction** to anticipate and preemptively address purchase barriers (e.g., offering financing to price-sensitive customers)
- **Applying personalization** to customize checkout flows based on customer preferences and history

- **Deploying automation** to resolve payment issues, apply discounts, and process transactions across complex backend systems

Key insight: Purchase success requires both functional excellence (smooth transactions) and psychological comfort (risk mitigation). AI prediction can identify likely abandonment points, personalization can adapt the experience to individual concerns, and automation ensures consistent execution at scale.

Post-Purchase: Building Relationships and Advocacy

After purchase, focus shifts from acquisition to retention and growth. This stage determines whether one-time buyers become loyal customers and eventually brand advocates. The post-purchase experience often matters more than the purchase itself in determining long-term relationships.

Satisfaction to Loyalty Progression:

- **Basic Satisfaction**: Products meet expectations; problems are resolved effectively
- **Preference Formation**: Positive experiences accumulate, creating habitual choice and reducing price sensitivity

- **Emotional Attachment**: Customers identify with brand values, choosing it even when alternatives seem attractive
- **Active Advocacy**: Customers voluntarily promote the brand, staking their reputation on recommendations

Modern loyalty transcends repeat purchasing. A streaming service that consistently surfaces relevant content builds confidence in future value. A retailer that remembers preferences while introducing welcome surprises balances reliability with delight. These experiences create switching costs that pure economic incentives cannot match.

Strategic Implications for Firms

Organizations must actively nurture post-purchase relationships through systematic engagement:

- **Managing expectations** through clear communication about delivery, setup, and usage
- **Monitoring satisfaction** proactively to identify and resolve issues before escalation
- **Creating value beyond the product** through education, community, and exclusive benefits
- **Recognizing and rewarding loyalty** in ways that feel personal rather than transactional
- **Using prediction** to anticipate customer needs

(replenishment, upgrades, complementary products) and identify churn risks
- **Leveraging personalization** to tailor ongoing communications, recommendations, and offers to individual usage patterns
- **Implementing automation** to handle routine service needs while escalating complex or emotional issues to humans

The customer journey provides a strategic framework for deploying AI capabilities. By understanding how customers progress from awareness through advocacy, organizations can identify where prediction, personalization, and automation create the most value. The goal isn't to automate the entire journey but to enhance it—using AI to eliminate friction, increase relevance, and enable human connections where they matter most.

AI's Core Capabilities: Personalization, Automation, and Prediction

While the customer journey provides the strategic framework for understanding customer interactions, artificial intelligence supplies the capabilities that transform how organizations execute at each stage. Three core AI capabilities have emerged

as fundamental to modern customer experience: personalization, automation, and prediction. These capabilities work both independently and synergistically—personalization tailors experiences to individual preferences and contexts, automation enables complex processes to run without human intervention while adapting to exceptions, and prediction anticipates future needs and behaviors before they fully manifest. What makes these capabilities truly revolutionary isn't just their individual power, but how they reinforce each other: prediction identifies what customers will want, personalization determines how to present it, and automation ensures consistent delivery at scale. The following sections examine each capability in depth, exploring not just what they do, but how they evolved from traditional approaches, how they technically function, and most critically, how they create value across different journey stages. Through detailed examples from companies at the forefront of AI implementation, you'll see how these capabilities move beyond incremental improvement to enable fundamentally new business models and customer relationships.

AI's Core Capabilities: Personalization

How does Spotify know exactly what songs to recommend in your Discover Weekly playlist? Why does TikTok's algorithm

seem to understand your interests better than your friends do? The answer lies in AI-powered personalization—a capability that has fundamentally transformed how organizations understand and serve individual customers throughout their journey.

Spotify Wrapped Now Includes an AI-Generated Podcast Analyzing Your Listening Habits

Want to hear two robots chat about your music taste? This Wrapped season, Spotify and Google teamed up to make that dream (or nightmare) a reality.



Understanding AI Personalization

Personalization through AI tailors content, products, and experiences to individual customers based on their unique characteristics, behaviors, and contexts. Think of it like having a friend who remembers every conversation, every preference you've mentioned, and every choice you've made—then uses

that knowledge to anticipate what you'll want next. Unlike traditional marketing that groups people into broad categories, AI creates a unique understanding of each individual.

Consider how Spotify works: It doesn't just know you like pop music; it understands you prefer upbeat pop on Monday mornings, indie pop while working, and nostalgic hits on Friday evenings. This isn't magic—it's pattern recognition applied to your listening behavior across thousands of data points.

Evolution from Traditional Approaches

Traditional personalization worked like a filing cabinet—customers went into folders labeled "young professionals" or "budget shoppers." If you bought running shoes, you'd see ads for running socks. Simple, but limited.

AI personalization works more like a personal assistant who learns your preferences over time. When Netflix recommends a show, it considers: what you've watched, when you stopped watching other shows, what time you typically watch, whether you're alone or with family, and patterns from millions of similar users. It discovers connections humans would never spot—like people who watch nature documentaries on Sunday mornings often enjoy slow-paced foreign films on weekday evenings.

The key difference: traditional systems follow rules humans

create ("if customer buys X, show Y"). AI systems discover their own patterns ("customers who do A, B, and C have an 85% chance of wanting D").

The Data Processing Revolution

This shift requires processing enormous amounts of information. While traditional systems might track 10 attributes per customer, AI systems analyze thousands: every click, scroll, pause, and purchase. Each interaction refines your profile, like an artist adding brushstrokes to an increasingly detailed portrait.

How Personalization Works

Personalization is an ongoing process where artificial intelligence continuously learns from your interactions to tailor experiences to your preferences. Every action you take—what you click, skip, search for, or linger on—provides data that helps the system understand your tastes. It looks for patterns across countless behaviors to infer when, how, and why you engage with certain types of content. Using these insights, AI generates responses that feel natural and relevant, whether that means suggesting a song that fits your mood, a product that matches your style, or a message timed to your habits. Each new interaction refines the system's understanding, allowing it to adapt and improve over time.

The result is a dynamic loop where learning and personalization happen continuously, making your experience increasingly accurate, timely, and engaging.

Dimensions of Personalization

Content Personalization adapts communication style. Netflix doesn't just recommend different shows to different people—it shows different artwork for the same show. Action fans see explosive scenes; romance viewers see intimate moments. Even descriptions adapt: technical viewers get plot details while casual viewers see emotional hooks.

Product Personalization goes beyond "customers also bought." Spotify's Discover Weekly combines your taste profile with exploration algorithms, introducing variety while maintaining relevance. It knows to suggest slightly adventurous choices on Friday (when you're more open) versus familiar comfort music on Monday mornings.

Experience Personalization modifies interfaces themselves. Notice how your Netflix homepage differs from others'? Some users see trending shows prominently; others see personalized categories like "Because you watched Stranger Things." These aren't random—they're based on learned interaction preferences.

Timing Personalization determines optimal engagement moments. The system learns you browse during lunch but purchase in evenings, check email at 8 AM but ignore

afternoon messages. This temporal intelligence prevents both missing opportunities and causing annoyance.

Modern Developments and Challenges

Generative AI Enhancement

Recent developments like ChatGPT are revolutionizing personalization. Instead of selecting from pre-written content, AI can now generate completely unique messages, product descriptions, or recommendations for each individual. Imagine product descriptions that adapt not just in content but in length, complexity, and style based on your preferences and expertise level.

The Context Challenge

Modern systems excel at contextual understanding. Netflix knows Sunday night viewing differs from Tuesday lunch breaks. Shopping for yourself differs from gift shopping. The same person exhibits different preferences across contexts—something traditional systems couldn't grasp.

Privacy and Consent Concerns

This deep personalization raises critical privacy questions throughout the customer journey. During awareness stages,

customers might not realize data collection has begun. By the purchase phase, extensive profiles exist. Post-purchase, data continues accumulating. Organizations must balance personalization benefits with transparency about data use, providing clear consent mechanisms and control options at each journey stage.

Filter Bubbles and Journey Implications

Over-personalization creates filter bubbles that affect the entire journey. In awareness stages, you might never discover products outside your established preferences. During consideration, you only see confirming information. This narrowing can lead to personalization fatigue—when customers feel trapped by their past choices, unable to explore or evolve.

Spotify addresses this through deliberate variety injection. "Release Radar" introduces new music within your comfort zone, while "Discover Weekly" pushes boundaries slightly further. This balance keeps the journey fresh while maintaining relevance.

Personalization Fatigue Across the Journey

Early in the journey, heavy personalization can feel creepy—"How did they know I was looking at that?" During consideration, excessive personalization might reduce the joy

of discovery. At purchase, over-personalized pricing can feel manipulative. Post-purchase, constant personalized upselling exhausts customers.

Implementation Considerations

Success requires quality data integration across all touchpoints. When systems remain disconnected—website behavior separate from store purchases—personalization fragments. Customers receive irrelevant recommendations or conflicting messages, disrupting their journey.

Organizations must also balance automation with human touch. While AI can personalize at scale, some journey moments require human connection. A customer complaint might need empathy no algorithm can provide. The key lies in using AI to enhance, not replace, human relationships.

Competitive Implications

Leaders in personalization show distinct advantages. Netflix retains subscribers through personalization worth $1 billion annually. Amazon generates 35% of revenue through recommendations. TikTok disrupted social media giants by learning preferences faster—understanding new users within hours rather than days.

Organizations struggling with personalization share common problems: disconnected data systems, treating

personalization as purely technical rather than strategic, and failing to balance personalization with discovery and privacy. They automate without understanding, personalize without purpose, and ultimately frustrate rather than delight customers.

AI's Core Capabilities: Automation

Why can customer service chatbots now solve complex problems that once required human agents? How do marketing campaigns optimize themselves across thousands of segments without human intervention? The answer lies in AI-powered automation—a capability that has transformed how organizations execute complex processes throughout the customer journey.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=302#oembed-1

Understanding AI Automation

AI automation enables systems to handle complex tasks with minimal human intervention while adapting to changing conditions. Think of the difference between a vending machine and a barista. The vending machine follows rigid rules: insert money, press button, receive drink. A barista understands context: they remember your usual order, suggest alternatives when ingredients run out, and adjust recommendations based on weather or time of day. AI automation brings that barista-like adaptability to digital systems at massive scale.

When you message Netflix customer service at 2 AM about a billing issue, you're likely chatting with AI that understands your problem, accesses your account history, and resolves issues that once required human agents. This isn't just following a script—it's understanding intent and taking appropriate action.

Evolution from Traditional Approaches

Traditional automation worked like a flowchart—if this, then that. When you called customer service and pressed "1" for billing, "2" for technical support, you navigated a rigid tree of pre-programmed responses. These systems failed when your problem didn't fit neat categories.

AI automation works like a skilled assistant who understands context and intent. When you type "my show keeps freezing," the system understands you're having streaming issues, checks your device type, internet speed, and recent viewing history, then provides personalized troubleshooting—or immediately escalates to human support if it detects frustration.

The key difference: traditional automation follows rules humans create, while AI automation learns patterns and adapts responses. It's the difference between a recipe (follow exact steps) and a experienced chef (adjust based on ingredients, taste, and situation).

The Data Processing Revolution

This shift enables processing unstructured information—the messy, human-generated content that comprises 80% of business data. Traditional systems needed perfect formatting. AI automation understands your hastily typed email, interprets voice messages with background noise, and even processes photos of damaged products for return requests.

How Automation Works Throughout the Journey

TAutomation functions as a continuous process in which AI systems interpret inputs, make decisions, act, and learn to

improve future performance. At every interaction, automation interprets what users mean rather than just what they say—for instance, distinguishing between someone who "can't find anything good" and someone whose "app won't load." It evaluates multiple factors such as customer history, context, and predicted outcomes to decide how to respond, whether by recommending content, offering a discount, or resolving a problem. Once a decision is made, the system coordinates actions across platforms, from adjusting search results to processing payments or managing returns. Each outcome then feeds back into the system's learning process. When a user engages positively with an automated response, the system strengthens that pattern; when they disengage, it adjusts accordingly. Over time, automation becomes more intelligent, responsive, and capable of managing complex customer interactions without direct human oversight.

Modern Developments in Automation

Generative AI Revolution

ChatGPT-style interfaces have transformed automation beyond simple responses. Instead of selecting from pre-written templates, these systems generate unique, contextual responses for each situation. Customer service bots now provide detailed, personalized explanations rather than generic

answers. They can explain complex billing situations, walk through technical troubleshooting, or even help with creative tasks like gift selection.

Spotify could theoretically use generative AI to create personalized playlist descriptions that explain why specific songs were chosen for you, written in a tone that matches your communication style. "Hey, noticed you've been into moody indie lately, so threw in that new Phoebe Bridgers track between your usual favorites."

Process Automation

Modern AI handles entire workflows. When you initiate a return on Amazon, automation orchestrates everything: generating shipping labels, scheduling pickups, processing refunds, updating inventory, and adjusting recommendations. This happens across dozens of systems, adapting to exceptions (international returns, damaged goods, warranty claims) without human intervention.

Communication Automation

AI now manages complete conversations. Unlike old chatbots that frustrated with "I don't understand," modern systems maintain context across long dialogues. They remember what you said ten messages ago, understand pronouns and references, and seamlessly hand off to humans when needed—briefing the agent so you don't repeat yourself.

The Human-Automation Balance in the Journey

When Automation Enhances the Journey

During **awareness and discovery**, automation excels at surfacing relevant content and answering basic questions instantly. Netflix's automated recommendations help you discover shows you'd never find manually.

In **consideration**, automation provides consistent information, comparison tools, and instant clarification. It never gets tired of answering the same questions and can serve thousands simultaneously.

At **purchase**, automation streamlines transactions, reduces friction, and ensures accuracy. One-click buying only works because automation handles the complexity behind scenes.

When Automation Hinders the Journey

During **emotional moments**, automation can feel cold. A customer upset about a deceased family member's account needs human empathy, not efficient processing. When someone's Netflix account was hacked and inappropriate content was watched, they need understanding, not just password reset instructions.

In **complex problem-solving**, automation struggles with nuanced situations. When your issue spans multiple problems

or requires creative solutions, human judgment becomes essential.

During **relationship-building moments**, automation cannot replace human connection. A loyal customer reaching a milestone or expressing gratitude deserves personal recognition.

Emotional Labor and Irreplaceable Human Connection

Some interactions require emotional labor—the work of managing feelings and expressions to fulfill emotional requirements of a job. When customers are frustrated, scared, or celebrating, they need genuine human connection. AI can simulate empathy but cannot truly understand loss, frustration, or joy.

Consider complaint handling: automation can process refunds and replacements efficiently, but an angry customer often needs to feel heard. The journey isn't just about solving problems—it's about feeling valued. Smart organizations use automation to handle logistics while humans handle emotions.

Implementation Considerations

Success requires comprehensive data integration. When systems remain disconnected, automation fails. Imagine Netflix's frustration if viewing history, billing, and recommendations operated separately—you'd get suggestions

for shows you've watched and bills for services you've cancelled.

Organizations must also determine automation boundaries. High-frequency, low-complexity tasks suit full automation. High-stakes, high-emotion situations need human involvement. The sweet spot often involves automation handling routine aspects while humans focus on exceptions and relationships.

Reliability becomes critical when automation handles customer journeys. Unlike traditional systems that clearly succeed or fail, AI automation can degrade subtly. A chatbot might still respond but give increasingly irrelevant answers. Organizations need sophisticated monitoring to detect not just failures but declining quality.

Competitive Implications

Leaders leverage automation strategically. Amazon's automation doesn't just reduce costs—it enables services impossible with human labor. Instant recommendations, one-click purchasing, and same-day delivery only work through extensive automation.

Google processes billions of ad auctions daily, each considering hundreds of variables in milliseconds. This speed and scale would require millions of human workers.

Organizations struggling with automation often automate wrong things. They digitize broken processes rather than

reimagining them. They view automation as cost-cutting rather than capability-building. While they automate simple tasks, competitors use automation to create entirely new customer experiences—24/7 instant support, predictive problem resolution, and personalized service at scale.

AI's Core Capabilities: Prediction

How does Amazon know you'll need printer ink before you realize you're running low? Why can credit card companies detect fraudulent transactions within milliseconds? The answer lies in AI-powered prediction—a capability that enables organizations to anticipate future events and behaviors throughout the customer journey.



Understanding AI Prediction

AI prediction forecasts future events by analyzing patterns in data. Think of it like how you predict your friend's behavior. You know Sarah orders dessert when she's stressed, Mike arrives 15 minutes late to everything, and Alex cancels plans when it rains. You learned these patterns through observation. AI does the same thing, but with millions of "friends" (customers) and thousands of behavioral signals.

When Netflix suggests you'll love a new show with 95% confidence, it's not guessing. It knows you binged similar shows, you watch sci-fi on weekends, you prefer series with

strong female leads, and thousands of people with your viewing patterns loved this show. It's pattern recognition at massive scale.

Evolution from Traditional Approaches

Traditional prediction worked like extending a ruler's edge—draw a line through past points and continue straight ahead. If you sold 100 umbrellas last January and 120 this January, traditional models predicted 140 next January. Simple, but wrong when unexpected factors emerge.

AI prediction works more like weather forecasting. It considers countless variables interacting in complex ways: temperature affects ice cream sales, but so do local events, social media trends, competitor promotions, and whether it's payday. The system discovers that ice cream sales spike when temperature exceeds 75°F, it's within three days of payday, and local schools are closed—patterns humans wouldn't connect.

The key difference: traditional methods assume tomorrow resembles yesterday. AI recognizes that small changes can have large effects, and that patterns themselves evolve.

The Data Processing Revolution

This shift requires processing enormous data variety. While traditional models might track sales and demographics, AI

prediction analyzes everything: how long customers hover over products, what they almost bought, when they shop, their social media sentiment, local weather, economic indicators. Each data point potentially improves predictions, but only if the system can separate signal from noise.

How Prediction Works in the Journey

*]:pointer-events-auto scroll-mt-[calc(var(--header-height)+min(200px,max(70px,20svh)))]" dir="auto" data-turn-id="request-68fee33d-631c-832b-b021-6e57269e3249-1" data-testid="conversation-turn-8" data-scroll-anchor="true" data-turn="assistant">

Prediction is a continuous learning process through which AI systems gather data, detect patterns, forecast outcomes, and refine their accuracy over time. As users move through different touchpoints, the system collects signals such as searches, clicks, comparisons, and usage behaviors to form a detailed picture of their preferences. Machine learning models then uncover subtle relationships within this data—connections too complex or numerous for humans to notice—revealing how context, timing, and habits shape decisions. Using these insights, the system generates forecasts about what users are likely to do next, from what content they'll enjoy to when they might disengage. Each prediction is then tested against real outcomes. When behavior aligns with

expectations, the model strengthens its understanding; when it doesn't, the system recalibrates. Over time, this continuous feedback loop allows predictions to grow more accurate, adaptive, and personalized to each individual's evolving patterns.

Dimensions of Prediction Across the Journey

Behavioral Prediction

This anticipates individual actions at each journey stage. During awareness, Spotify predicts which new artists might interest you. During consideration, it predicts whether you'll upgrade to premium. During usage, it predicts when you'll listen (workout mornings, commute evenings) and what mood you're in (energetic Mondays, mellow Sundays).

Demand Prediction

Organizations forecast needs before customers realize them. Amazon's "Subscribe & Save" doesn't wait for you to reorder—it predicts when you'll run out based on purchase frequency, household size, and usage patterns. This shifts the journey from reactive (customer realizes need) to proactive (company anticipates need).

Risk Prediction

Systems identify problems before they impact the journey. Credit card companies analyze every transaction in milliseconds, comparing against your personal patterns. That coffee purchase seems normal, but not at 3 AM in a city you've never visited. The system balances security (stopping fraud) with convenience (not blocking legitimate purchases).

Lifetime Value Prediction

Early journey signals predict long-term relationships. Netflix knows within your first week whether you'll become a loyal subscriber based on viewing diversity, session length, and engagement patterns. This shapes everything from content recommendations to retention offers, customizing the journey based on predicted value.

Modern Developments and Challenges

Generative AI Enhancement

New AI capabilities are revolutionizing prediction. Rather than just forecasting what you'll want, systems can now generate personalized content explaining why you'll like something. Imagine Spotify creating custom podcast segments discussing why certain songs match your mood, or Netflix

generating personalized preview trailers emphasizing aspects you care about.

Prediction Confidence and Journey Design

Predictions come with confidence levels that should influence journey design. When Netflix is 95% confident you'll love a show, it features it prominently. At 60% confidence, it might include it lower in recommendations. At 30%, it might only show it if you're actively browsing.

This uncertainty affects journey touchpoints differently. High-confidence predictions enable bold moves (Amazon's anticipatory shipping). Low-confidence predictions require hedging (showing multiple options). Organizations must design journeys that gracefully handle prediction uncertainty.

Prediction Horizons and Journey Planning

Different timeframes require different approaches:

Immediate predictions (next few minutes) guide real-time experiences. Netflix predicts whether you'll continue watching or browse for something new, adjusting interface accordingly.

Short-term predictions (days to weeks) inform engagement strategies. Spotify predicts when you'll need new music, timing "Discover Weekly" accordingly.

Long-term predictions (months to years) shape strategic

relationships. Subscription services predict lifetime value, determining investment in acquisition and retention.

Each horizon uses different data and affects different journey decisions, but they must work together coherently.

The Feedback Loop Challenge

Here's a simple way to understand feedback loops: predictions can create the outcomes they predict. Imagine Netflix predicts you won't like romantic comedies, so it stops showing them. You never watch any (because you don't see them), confirming the prediction. But was the prediction right, or did it create its own truth?

This self-fulfilling prophecy affects entire journeys. If Amazon predicts you're a bargain shopper and only shows sale items, you might never discover premium products you'd actually prefer. Smart systems deliberately test their assumptions, occasionally showing content outside predictions to prevent these loops.

Implementation Considerations

Success requires quality data and careful bias management. If historical data reflects past problems—like discrimination in lending—predictions perpetuate them. Organizations must audit data for representativeness and fairness.

The explainability challenge grows as predictions influence

important decisions. While "black box" predictions work fine for song recommendations, credit decisions require explanation. Organizations balance accuracy with transparency based on context.

Competitive Implications

Leaders use prediction strategically throughout journeys. Amazon doesn't just predict purchases—it reshapes the entire journey around anticipation. Stitch Fix built its business on prediction, sending clothes before you know you want them.

Organizations struggling with prediction remain reactive. While Amazon ships products before orders, traditional retailers still struggle with inventory. While Netflix retains subscribers through predictive personalization, traditional media loses audiences they never saw leaving.

Conclusion

The customer journey and AI capabilities represent two sides of the same strategic imperative: creating superior customer experiences in an increasingly digital world. The journey framework reveals where value is created or destroyed—from the moment customers recognize needs through their evolution into brand advocates. AI capabilities show how organizations can transform these moments, using prediction

to anticipate needs, personalization to tailor responses, and automation to execute flawlessly at scale. Yet the most important lesson from this chapter isn't about individual technologies or journey stages—it's about integration and purpose. Organizations that treat AI as merely a cost-cutting tool or implement capabilities without understanding the customer journey will achieve incremental improvements at best. In contrast, those that thoughtfully orchestrate AI capabilities across the entire journey create compound advantages: customers feel understood rather than analyzed, supported rather than processed, and valued rather than monetized. As you've seen through examples from Netflix, Amazon, Spotify, and others, success requires both strategic clarity about the customer journey and technical sophistication in deploying AI capabilities. The organizations that master this combination don't just serve customers better—they fundamentally redefine what customers expect from every other organization they encounter. Understanding these frameworks isn't just about comprehending current practice; it's about preparing for a future where the integration of human insight and artificial intelligence becomes the primary source of competitive advantage.

Key Takeaways

The Customer Journey Framework

The **Customer Journey** is the complete map of a customer's interactions with a brand, from initial need recognition to long-term loyalty. It provides a strategic roadmap for identifying where and how to create value. The four primary stages are:

- **Awareness:** The stage where a customer recognizes a need or problem. The key for organizations is to help customers articulate their problems through educational content, not aggressive selling.
- **Consideration:** The active research phase where customers evaluate and compare different solutions. Organizations must simplify this complex process by providing transparent information and decision aids.
- **Purchase:** The critical moment of transaction. The goal is to minimize **friction** (obstacles to completion) and maximize confidence through seamless processes and trust signals like guarantees.
- **Post-Purchase:** The phase after the sale that determines long-term loyalty. The focus shifts from acquisition to building a relationship through proactive support, continued value, and engagement, turning satisfied customers into brand advocates.

AI's Core Capabilities

AI provides three transformative capabilities that revolutionize how organizations engage customers at each stage of their journey. They work together to anticipate needs, tailor interactions, and deliver experiences at scale.

- **Personalization:** This is the AI-powered tailoring of content, products, and experiences to an individual, not a broad segment. It moves beyond simple, rule-based targeting (e.g., "customers who bought X also bought Y") to discover complex patterns in user data and create a unique experience for each person, like Netflix's recommendation engine.
- **Automation:** This involves using AI to handle complex tasks with adaptability and minimal human intervention. Unlike traditional, rigid automation (e.g., a simple phone menu), AI automation can understand intent, manage entire workflows, and adapt to changing conditions, like a smart chatbot that resolves a complex billing issue.
- **Prediction:** This is the ability to forecast future behaviors and events by analyzing patterns in vast amounts of data. It moves beyond simple linear forecasting to a more sophisticated, weather-forecast-like model that considers countless interacting variables to anticipate customer needs, churn risk, or

potential fraud before they happen.

5.

AI AND COMPETITIVE AND CUSTOMER INTELLIGENCE

AI and customer and competitive intelligence

This chapter is based on Jayawardena, Behl, Thaichon, and Quach (2023),

Market orientation is a fundamental concept in marketing. It represents a shift from a product-centric approach to a market-centric one, emphasizing the importance of understanding and responding to customer needs and competitors' actions. A market-oriented organization continuously collects information about customers, competitors, and market trends, using this intelligence to create superior value for customers.

Customer orientation refers to the gathering, analyzing, and interpreting information about customers to gain insights into their behaviors, preferences, and needs. It encompasses a wide

range of data, from demographic information to purchase history and interactions with the brand. The goal of customer intelligence is to develop a deep understanding of customers, enabling businesses to personalize their offerings, improve customer experiences, and build lasting relationships.

Competitive orientation involves the systematic collection and analysis of information about competitors and the overall competitive landscape. This includes monitoring competitor strategies, products, pricing, and market positioning. Competitive intelligence helps businesses identify threats and opportunities in the market, benchmark their performance, and make informed strategic decisions.

Together, customer and competitive intelligence form the foundation of market orientation, providing the insights necessary for businesses to align their strategies with market realities and customer expectations.

The challenge facing modern marketing organizations is not information scarcity but overwhelming abundance. Customer interactions generate digital traces across websites, mobile applications, social media platforms, email systems, customer service channels, and physical retail environments. Simultaneously, competitive actions—product launches, pricing adjustments, promotional campaigns, strategic partnerships—occur with increasing frequency and visibility across multiple digital channels. Organizations possess more data than ever before, yet many struggle to convert this information into actionable insight due to fragmentation

across systems, velocity of change, and sheer scale of available information.

Artificial intelligence offers potential resolution to this information paradox. In marketing contexts, AI encompasses computational systems that perform tasks traditionally requiring human cognitive capabilities: pattern recognition across large datasets, prediction of future outcomes based on historical patterns, natural language processing and generation, and automated decision-making within defined parameters. These capabilities enable marketers to process information at scales previously unattainable, identifying patterns and generating insights that would remain invisible through manual analysis alone.

This chapter examines how AI enhances marketing information systems across two domains essential to market orientation. We explore how organizations understand customers and competitors.

Research Methods Foundation

Before examining how AI enhances marketing information systems, we must understand the research processes that form their foundation. Marketers gather information through diverse methods, each suited to particular questions and contexts. These methods apply whether organizations are

conducting one-time studies to answer specific questions or continuously monitoring their markets for strategic awareness.

Information needs define what organizations must gather. About customers, marketers need to understand who buys their products, what customers value, how they make purchase decisions, why they choose particular brands, which customers are most profitable, and what unmet needs exist that products could address. About competitors, marketers need to understand rival strategies and capabilities, competitive product features and pricing, how customers perceive alternatives, what industry dynamics shape competition, and where strategic gaps create opportunities. These questions determine appropriate research methods and data sources.

Qualitative methods seek rich understanding of experiences, motivations, and contexts. In-depth interviews involve extended one-on-one conversations with customers exploring their experiences, decision processes, and needs. Focus groups bring together small groups to discuss products, brands, or concepts, revealing not just individual perspectives but how people influence each other's thinking. Observational research and ethnography involve watching customers in natural settings—their homes, workplaces, or shopping environments—to understand actual behaviors rather than self-reported accounts. When **Procter & Gamble wanted to understand global laundry practices**, ethnographic researchers observed families in their homes across multiple

countries, revealing that Indian women often washed clothes in buckets rather than machines while Japanese consumers valued concentrated detergents due to limited storage space. These insights emerged from observation rather than surveys, illustrating how qualitative research reveals the "why" behind customer behavior.

Quantitative methods provide numerical data amenable to statistical analysis. Surveys systematically collect information from many respondents, measuring attitudes, preferences, behaviors, and demographics through structured questionnaires. Experiments test cause-and-effect relationships through controlled manipulation of variables, such as testing whether different pricing levels affect purchase rates. Behavioral tracking and transactional data reveal what customers actually do through website analytics, mobile app usage, point-of-sale systems, and loyalty programs. This revealed preference data shows actual behavior under real conditions, complementing attitudinal research that captures what customers think and feel.

The distinction between primary and secondary data shapes research approaches. Primary data is information collected specifically for current questions through methods the researcher controls—surveys sent to target customers, focus groups moderated by the research team, or experiments conducted in controlled settings. Secondary data already exists, collected previously for other purposes but potentially relevant to current questions. Industry reports, government statistics,

academic research, competitor financial disclosures, and internal company records all constitute secondary data sources. Secondary research typically costs less and proceeds faster than primary research but may not perfectly match current information needs.

Research quality depends critically on sampling and measurement decisions. When studying populations too large to survey exhaustively, researchers select samples designed to represent the broader group. Probability sampling methods use random selection to ensure every population member has a known chance of inclusion, enabling statistical generalization from sample to population. Non-probability methods sacrifice generalizability for accessibility or theoretical relevance, appropriate when representative samples prove impossible or when depth matters more than breadth.

Measurement quality hinges on questionnaire design for surveys or interview protocol development for qualitative studies. Questions must be clear, unbiased, and appropriate to respondents' knowledge and willingness to answer. Leading questions, double-barreled questions that ask multiple things simultaneously, or questions using terminology respondents don't understand all compromise data quality. A question like "How satisfied are you with our excellent customer service?" introduces bias by embedding a positive evaluation. Better formulation asks "How satisfied are you with the customer service you received?" without presuming its quality.

These research fundamentals establish the foundation upon

which AI tools build. Understanding methodological principles enables marketers to recognize that AI accelerates certain activities but does not eliminate the need for sound research design, appropriate methodology selection, or human interpretation of findings. Technology amplifies good research practice but cannot rescue poorly conceived studies.

Understanding Customers Through Data and AI

Customer information represents systematic gathering and analysis of data about customers to inform marketing strategy and execution. Organizations need comprehensive understanding of customer behaviors, preferences, and needs to make effective decisions about segmentation, targeting, positioning, product development, and relationship management.

What We Need to Know About Customers

Customer understanding encompasses multiple information types that together create comprehensive knowledge. Behavioral information describes what customers actually do—their purchase patterns, channel preferences, product

usage, decision-making journeys, and response to marketing stimuli. This revealed preference data shows how customers behave under real conditions but doesn't explain motivations. Attitudinal information captures what customers think and feel through satisfaction surveys, brand perception studies, and voice-of-customer programs. It reveals preferences, perceptions, and stated intentions but may diverge from actual behavior when social desirability or intention-action gaps emerge. Needs-based information identifies underlying problems customers are trying to solve, often using frameworks like jobs-to-be-done that focus on customer goals rather than product features.

Effective customer understanding also requires segmentation—identifying groups of customers who warrant distinct treatment based on shared characteristics, behaviors, or needs. Traditional segmentation uses demographics, geography, or broad psychographic categories. More sophisticated approaches segment by behavior patterns, needs, or profitability, creating actionable groups that inform targeting and positioning decisions.

Traditional Customer Information Methods

Organizations employ diverse methods to generate customer understanding. Transaction analysis examines purchase histories to identify patterns: which products are bought

together, seasonal fluctuations, migration as customers adopt additional products over time, and differences between high-value and low-value customers. RFM segmentation groups customers by Recency of last purchase, Frequency of purchases, and Monetary value to identify high-value customers deserving different treatment than occasional buyers.

Surveys and satisfaction tracking systematically measure customer attitudes, preferences, and experiences. Net Promoter Score studies track willingness to recommend, providing a metric for loyalty trends. Brand perception research measures awareness, associations, and positioning relative to competitors. These quantitative approaches enable tracking over time and comparison across segments.

Customer interviews and voice-of-customer programs gather qualitative feedback through structured conversations, capturing rich detail about experiences, pain points, and unmet needs. **LEGO's decision to develop the Friends line** specifically for girls emerged from qualitative research revealing that girls wanted different building experiences than boys, not that they wouldn't build at all. This attitudinal insight challenged internal assumptions and informed a successful strategic shift.

Observational research and ethnography involve watching customers in natural settings to understand actual behaviors. GE's face-to-face customer engagement initiative sent executives into customer environments to observe problems

firsthand, generating qualitative insights about operational challenges that quantitative satisfaction scores never revealed. These observations then informed quantitative studies measuring the prevalence of identified issues across the customer base, combining qualitative depth with quantitative scale.

Customer journey mapping visualizes how customers move through awareness, consideration, purchase, and post-purchase stages across touchpoints, identifying pain points where interventions could improve experience or moments of truth that disproportionately influence satisfaction.

Traditional methods face limitations related to scale, speed, and fragmentation. Manual analysis of transaction data limits researchers to summary statistics rather than identifying nuanced micro-patterns. Qualitative research faces a breadth-depth tradeoff where researchers can conduct deep interviews with small samples or surface-level analysis of larger datasets, but analyzing hundreds of lengthy interviews proves prohibitively time-consuming. Customer data typically fragments across separate systems—website analytics, CRM databases, point-of-sale systems, call center records—making it difficult to construct complete customer views. These limitations create opportunities for AI enhancement.

How AI Enhances Customer Understanding

Artificial intelligence amplifies customer understanding capabilities in ways that address traditional limitations while introducing new considerations. AI applications span the entire information gathering process from design through collection, analysis, and insight generation.

Survey Design and Optimization

AI tools now assist with questionnaire development, addressing one of research's most challenging phases. Natural language generation models suggest question formats, refine ambiguous phrasing, and identify potential bias in wording. Rather than starting from scratch, researchers can prompt AI with research objectives—testing brand recall, measuring purchase intent, evaluating feature preferences—and receive multiple draft questions within seconds. These tools also optimize survey flow and logic, designing branching patterns that minimize respondent fatigue while maximizing data quality.

Adaptive surveys use AI-powered conversational interfaces that adapt questioning based on previous responses. Unlike traditional surveys that ask identical questions to all respondents regardless of relevance, conversational surveys clarify confusing items and skip irrelevant sections. A

customer indicating they haven't purchased in six months need not answer detailed questions about recent purchase satisfaction. This adaptive approach improves completion rates while generating richer, more targeted data.

AI systems also enable automated data quality monitoring, detecting fraudulent survey responses in real-time. They identify patterns like straight-lining where respondents select the same response repeatedly, impossibly fast completion times, nonsensical open-ended responses, or IP addresses associated with survey farms. This fraud detection improves data quality by removing problematic responses before they contaminate analysis.

For example, a **fintech research team used AI to test different question orderings** for a mobile survey, reordering questions to prioritize emotionally engaging content early in the flow, resulting in a 22% increase in completion rates

Behavioral Analysis at Scale

Machine learning transforms behavioral intelligence through its ability to process millions of customer journeys and identify micro-patterns invisible in aggregated reports. Traditional analysis might segment customers into broad categories like "frequent buyers" and "occasional buyers." AI identifies thousands of behavioral micro-segments defined by nuanced patterns: customers who browse extensively but purchase infrequently, customers who purchase only during

promotional periods, customers whose basket sizes steadily increase over time, or customers whose category mix suggests they're consolidating purchases from multiple retailers.

Netflix exemplifies behavioral analysis at scale. Rather than segmenting viewers into handful of categories like "likes comedy" or "likes drama," their recommendation system identifies **thousands of taste clusters** based on viewing patterns, rating behaviors, browsing without watching, and completion rates. This granular segmentation enables hyper-targeted content recommendations that would be impossible with manual analysis.

Real-time behavioral processing enables immediate response to customer actions. hen **Starbucks' mobile app** detects a customer near a store location they haven't visited previously, it can trigger a location-specific welcome offer within seconds. When an e-commerce system detects abandonment patterns suggesting the customer is confused or comparison shopping, it can intervene with chat support or time-limited incentives. This real-time capability transforms behavioral data from retrospective analysis tool to dynamic engagement enabler.

Integration Across Touchpoints

Customer Data Platforms use machine learning for identity resolution—determining which activities across devices, browsers, email addresses, and physical locations belong to the

same individual. This solves a fundamental data fragmentation problem: a customer might browse your website on a laptop, research products on a mobile phone, call customer service with questions, then purchase in a physical store. Without identity resolution, these appear as separate individuals rather than one customer's complete journey.

Starbucks demonstrates effective integration, connecting mobile app usage, in-store purchases via loyalty cards, website interactions, and customer service contacts to understand each customer's preferences and behaviors comprehensively. This unified view reveals complete customer journeys that remain invisible when data stays siloed in separate systems. The platform can then determine that customers who use the mobile app to order ahead have significantly higher visit frequency than those who don't, informing strategies to drive app adoption.

However, identity resolution introduces privacy considerations. Connecting activities across contexts enables comprehensive customer understanding but also creates detailed behavioral profiles that customers may find intrusive. Organizations must balance analytical value against privacy concerns and regulatory requirements like GDPR and CCPA that govern how personal data can be collected, connected, and used.

Predictive Customer Analytics

Rather than only describing current customer states, AI forecasts future behaviors and outcomes. Churn prediction models analyze behavioral signals—login frequency, feature usage, support ticket volume, payment delays, competitor website visits—to assign each customer a probability of defecting within defined timeframes. Organizations can then deploy retention interventions toward high-risk valuable customers before they actually leave. When models identify a previously active user whose engagement has declined sharply, customer success teams can proactively reach out with assistance, special offers, or product education before the customer cancels. For example, **Spotify uses predictive models** to analyze users' listening habits and predict when engagement patterns indicate potential churn.

Customer lifetime value forecasting attempts to predict total profit a customer will generate over their entire relationship with the firm. This requires estimating future purchase frequency, transaction values, retention duration, and associated costs—each uncertain. Airlines use lifetime value predictions to determine which frequent fliers merit elite status and which should receive targeted incentives. Small errors in component estimates compound, producing large lifetime value estimation errors, particularly for customers with long expected relationships. Organizations may therefore

allocate acquisition spending or service resources based on unreliable projections.

Next-best-action recommendations use predictive models to suggest optimal customer engagement strategies. Based on individual customer profiles and past responses to different approaches, systems recommend which product to offer, which message to send, or which channel to use for maximum response probability. A bank's system might determine that one customer responds well to email offers while another never opens marketing emails but does respond to in-app messages, tailoring outreach accordingly. For example, **Bank of America's virtual assistant Erica** serves nearly 50 million users with over 3 billion interactions, providing more than 1.7 billion tailored insights to assist clients in managing spending, savings, and investments, with AI determining optimal engagement strategies for each customer

Text and Sentiment Analysis

Natural language processing addresses the traditional breadth-depth tradeoff in qualitative research. While manual coding of interview transcripts or open-ended survey responses limits researchers to relatively small samples, NLP enables thematic analysis and sentiment assessment of thousands of text responses.

Sentiment analysis automatically assesses emotional tone across customer reviews, social media comments, survey

responses, and customer service transcripts, classifying content as positive, negative, neutral, or mixed. An electronics retailer receiving 50,000 product reviews monthly can use NLP to identify which products receive disproportionately negative sentiment, which specific features drive dissatisfaction, and how sentiment trends over time or varies by customer segment. **Amazon uses sentiment analysis** with AI and natural language processing to detect how customers feel about products based on their written reviews, analyzing 750 million customer reviews to classify sentiment and highlight what's driving emotions like price, delivery, or product quality.

Thematic coding categorizes responses into recurring themes and sub-themes. A hotel chain analyzing open-ended guest feedback might discover through NLP that "cleanliness" mentions cluster into distinct themes: bathroom cleanliness, bed linens, public areas, and dining facilities. This granularity reveals that poor sentiment associates specifically with bathroom cleanliness rather than cleanliness generally, focusing improvement efforts appropriately. **Marriott uses AI-powered sentiment analysis tools like Chatmeter** to monitor and analyze guest reviews across multiple platforms in real-time, enabling immediate service recovery and personalized responses, with the company tracking guest emotions expressed through online reviews and social media.

However, NLP systems struggle with nuance that humans readily understand. Sarcasm often reverses apparent sentiment: "Oh great, another software update that breaks

everything" is negative despite containing "great." Mixed sentiment statements like "The product works well but costs too much" require understanding that overall assessment depends on whether performance or price matters more to that customer. Cultural context and domain-specific terminology also challenge automated systems. Researchers must validate NLP outputs against human coding of sample data to ensure acceptable accuracy before trusting large-scale analysis.

Segmentation Through Clustering

Unsupervised machine learning algorithms identify customer segments from behavioral and attitudinal data without requiring researchers to specify segment characteristics in advance. Unlike traditional segmentation where analysts define segments based on demographics or suspected behavioral differences, clustering algorithms find naturally occurring groups based on similarity across many variables simultaneously.

An online retailer might use clustering on variables including purchase frequency, average order value, product category preferences, price sensitivity, promotional response, mobile versus desktop usage, and time between visits. The algorithm might identify segments like: price-sensitive frequent browsers who purchase only during sales, high-value convenience-focused customers who purchase full-price across

many categories, mobile-first shoppers who make frequent small purchases, and infrequent high-value purchasers who buy primarily seasonal gifts. These data-driven segments often reveal customer groups that demographic segmentation misses. **Sephora segments Beauty Insider members** into behavioral clusters beyond simple spending tiers, identifying groups like "Champions," "At-Risk Customers," and "Potential Loyalists," with targeted strategies for each segment including personalized email campaigns and re-engagement offers

However, data-driven segmentation requires careful validation. Algorithms find statistically distinct groups, but these may lack strategic meaning. A segment defined by "customers who shop on Tuesday mornings" might be statistically real but strategically meaningless if nothing about Tuesday morning shopping suggests different needs or profitability. Marketers must evaluate whether segments are actionable—can you describe each segment meaningfully, take different actions for each, and reach them through available channels? Segments that fail these tests provide statistical curiosity without practical value.

Synthetic Respondents

Among the most controversial emerging customer research applications are synthetic respondents—AI-generated personas designed to simulate human responses in research

scenarios. Large language models trained on vast text corpora can generate responses to survey questions or interview prompts that approximate how specific demographic or psychographic segments might answer.

Proponents argue this enables rapid, low-cost exploratory research for preliminary concept testing, reaching hard-to-access populations, or iterative refinement before investing in studies with human participants. A pharmaceutical company might use synthetic patient personas to map rare disease patient journeys before engaging scarce real patients, allowing researchers to develop more informed protocols for subsequent human research. A consumer goods company might test dozens of concept variations with synthetic respondents to identify the most promising options before expensive testing with actual consumers.

Current evidence suggests both potential and serious limitations. Researchers at the Wisconsin School of Business found that synthetic respondents could provide directionally useful qualitative feedback for early-stage concept exploration. Some market research practitioners predict that synthetic responses may comprise the majority of research data within three years as the technology matures.

However, synthetic respondents present methodological concerns warranting extreme caution. First, AI models generate outputs based on statistical patterns in training data, not genuine lived experience, emotions, or consciousness. They cannot provide authentic insight into human

motivation, which often involves subconscious processes, contextual factors, and individual experiences that statistical models cannot fully capture. A synthetic respondent can state what pattern in training data suggests a demographic group might prefer, but cannot experience the emotional response or contextual factors that actually drive human preference.

Second, training data biases propagate to synthetic outputs. If training corpora over-represent certain demographic groups or perspectives, synthetic respondents will reproduce these biases, potentially misrepresenting underrepresented populations. A synthetic persona meant to represent elderly consumers might reflect younger people's stereotypes about elderly consumers rather than authentic elderly perspectives if training data contains more content written about elderly people than by them.

Third, AI models sometimes generate plausible-sounding but inaccurate responses, a phenomenon termed "hallucination." Synthetic respondents can confidently assert facts or preferences having no basis in reality, misleading researchers who cannot easily distinguish authentic patterns from artifacts. Unlike human respondents whose statements, while sometimes inaccurate, reflect some actual experience or belief, synthetic statements may reflect nothing beyond statistical word associations.

Fourth, synthetic respondents cannot capture emerging trends or preferences not represented in historical training data. If customer needs are shifting due to new technologies,

economic conditions, or cultural changes, synthetic respondents trained on past data will miss these developments. They are retrospective by nature, making them inappropriate for research attempting to identify novel opportunities or changing preferences.

Organizations considering synthetic respondents should apply stringent constraints. They are inappropriate for decisions requiring understanding of authentic human experience, for studying populations underrepresented in AI training data, for sensitive topics where misrepresentation could cause harm, and for final validation before product launch or campaign deployment. If used at all, they should be limited to early-stage exploratory research where directional feedback suffices and error tolerance is high, always followed by validation with human participants. The emerging consensus positions synthetic respondents as tools for hypothesis generation rather than replacements for human research, though even this limited role remains controversial among research methodologists.

Critical Assessment and Limitations

AI customer information systems amplify analytical capabilities while introducing new constraints and failure modes. Data quality determines AI performance more than algorithm sophistication. Systems trained on biased,

incomplete, or outdated data produce unreliable outputs regardless of technical sophistication. Organizations often underestimate the data cleaning, integration, and maintenance work required for effective AI implementation.

Privacy and ethical considerations intensify as customer information systems become more comprehensive and personalized. Detailed behavioral tracking enables better customer understanding but also creates surveillance concerns. Algorithmic personalization can devolve into manipulation when systems optimize for short-term revenue extraction rather than long-term customer value. Organizations must balance analytical capability against customer trust and regulatory compliance.

Human judgment remains essential for strategic interpretation. AI identifies patterns but cannot explain why patterns exist or evaluate strategic significance. A system might detect that customers who purchase product A frequently also purchase product B, but humans must determine whether this reflects complementary needs worth bundling together, sequential purchase patterns suggesting lifecycle stages, or mere coincidence. Connecting patterns to actionable strategy requires business judgment that algorithms cannot replicate.

Understanding Competitors and Markets Through Data and AI

Competitive information involves systematic collection and analysis of data about competitors and the overall competitive landscape to inform strategic decisions. While customer information looks inward at an organization's own customers, competitive information looks outward at rivals, potential entrants, substitute products, and industry dynamics that shape competitive success. Effective competitive information enables organizations to anticipate competitor moves, identify strategic gaps and opportunities, benchmark performance, and position offerings advantageously.

What We Need to Know About Competitors

Competitive understanding encompasses multiple information needs. Organizations need to understand competitor strategies—which markets they prioritize, which customer segments they target, how they position their offerings, and what capabilities they are building for future competition. They need visibility into competitor product features, service levels, pricing strategies, and promotional

activities to inform their own offering development and positioning decisions.

Understanding how customers perceive competitive offerings provides essential context for positioning strategy. Competitors' self-description matters less than customer perceptions of their strengths, weaknesses, and differentiation. Organizations also need awareness of industry dynamics—consolidation trends, technological changes, regulatory developments, new entrant threats, and substitute products that could disrupt existing competitive patterns.

Competitive information ultimately reveals strategic gaps—unmet customer needs, underserved segments, or capability combinations that no competitor currently offers—creating opportunities for differentiation and advantage.

Traditional Competitive Information Methods

Organizations gather competitive information through diverse sources and methods. Secondary sources include publicly available information that competitors disclose. Financial reports reveal revenue trends, profitability, investments, and strategic priorities for public companies. Press releases and earnings calls provide management perspective on strategy and performance. Patent filings reveal R&D directions years before product launches. Job postings indicate capability building

and geographic expansion plans. Website content, marketing materials, and social media activity show positioning, messaging, and market focus. News coverage and industry analyst reports provide external perspectives on competitor strategies and performance.

Primary competitive information gathering involves methods requiring direct effort. Win-loss interviews where sales teams systematically debrief customers who chose competitors reveal authentic customer comparison criteria and competitor advantages from the buyer perspective. Mystery shopping enables experiencing competitor offerings firsthand—their service quality, sales process, product presentation, and pricing. Supplier and distributor conversations sometimes reveal competitor approaches, order volumes, or strategic shifts, though ethical boundaries prohibit misrepresentation or theft of confidential information. Trade show intelligence gathered through observation and conversations with competitor representatives provides visibility into product development, messaging, and market positioning.

Analytical frameworks structure competitive information into strategic insight. Competitor profiling creates comprehensive assessments of individual rivals covering their strategy, target markets, capabilities, strengths and weaknesses, financial resources, and likely future moves. Strategic group mapping identifies clusters of competitors pursuing similar strategies, revealing positioning space and mobility barriers

between groups. Industry structure analysis applies Porter's Five Forces framework to assess threats from existing rivals, potential entrants, substitute products, supplier power, and buyer power, illuminating how industry structure shapes profitability and strategic options.

SWOT analysis translates competitive information into strategic implications for the focal organization, identifying internal strengths and weaknesses relative to competitors while highlighting external opportunities and threats revealed through environmental scanning. War gaming and scenario planning use competitive information to anticipate how rivals might respond to strategic moves, enabling organizations to develop contingency plans before committing resources.

Traditional competitive information faces limitations. Manual monitoring of multiple competitors across diverse sources proves time-consuming and often incomplete. Weekly website checks miss mid-week changes. Analysts can track only limited numbers of competitors systematically. Published information arrives with time lags—financial reports describe past quarters, not current conditions. Information overload occurs when organizations accumulate competitive data without analytical frameworks that translate information into strategic insight.

How AI Enhances Competitor Understanding

Artificial intelligence amplifies competitor understanding capabilities in ways that address traditional limitations while introducing new considerations.

Competitive Pricing and Promotion Tracking

AI-powered pricing monitoring tracks competitor prices, discounts, and promotional timing across thousands of products in near real-time. Rather than manually checking competitor websites weekly, automated systems detect price changes within hours, identify promotional patterns, and alert managers to significant competitive moves.

Harvard Business Review's guide to real-time pricing emphasizes that advanced AI models consider much more than what competitors are charging—they analyze product availability, demand patterns, and consumer behavior to tailor pricing responses. Traditional approaches where retailers simply charge "X percent less than the lowest-price competitor" miss significant opportunities because they fail to account for these contextual factors.

Amazon and Walmart exemplify competitive price monitoring at massive scale. Both retailers use AI-powered systems to track each other's pricing thousands of times daily,

with Amazon's dynamic pricing algorithms adjusting product prices in real-time based on competitor pricing movements. According to Harvard Business School's analysis of dynamic pricing, platforms like Expedia employ similar systems to collect real-time airline data and calculate prices that factor in market demand, competitors' pricing, and time remaining before departure.

Machine learning identifies recurring pricing patterns that inform strategic responses. A retailer might discover that a key competitor consistently reduces prices every Friday afternoon to drive weekend traffic, or launches promotional campaigns on the first Monday of each month. These insights enable anticipating competitive moves rather than merely reacting to them. However, over-reliance on automated competitive pricing creates risks—blindly matching competitor prices can trigger destructive price wars that erode profitability for all players. Organizations must balance competitive responsiveness with strategic pricing discipline.

Competitor Review and Sentiment Analysis

Natural language processing analyzes thousands of competitor customer reviews to identify their strengths and weaknesses from the customer perspective. Rather than reading competitor reviews manually, NLP systems process reviews

at scale to identify frequently mentioned themes, associated sentiment, and how customer perceptions change over time.

This competitive intelligence reveals specific vulnerabilities to exploit or strengths requiring response. If analysis of 15,000 hotel reviews reveals that a competitor's customers consistently complain about bathroom cleanliness in 28% of reviews while only 4% mention restaurant quality, this intelligence directly informs both marketing positioning and operational priorities. Your marketing messages can emphasize superior cleanliness standards, while operations teams ensure similar weaknesses don't exist in your properties.

Research on AI-powered sentiment analysis shows that companies using these systems see an average 25% increase in customer satisfaction by identifying and addressing pain points revealed in competitor reviews. The analysis goes beyond simple positive or negative categorization to capture emotional nuances through natural language processing.

Miro's competitive analysis framework demonstrates how AI monitors reviews, testimonials, and social mentions to understand competitor strengths and pain points from the customer perspective, enabling teams to identify which content types drive the most engagement and spot emerging conversation topics in their industry.

However, NLP struggles with the same nuances discussed in customer sentiment analysis—sarcasm, cultural context, and mixed sentiment. A review stating "Great, another overpriced hotel charging premium rates for mediocre service"

appears positive due to "great" and "premium" despite being deeply negative. Competitor sentiment analysis also introduces ethical boundaries—while analyzing public reviews is standard practice, some organizations cross into questionable territory by mining competitor customer service interactions or private communities. Organizations must ensure competitive intelligence gathering remains ethical and legal.

Competitive Benchmarking and Gap Analysis

AI systems continuously track competitor product capabilities, features, messaging, and positioning to maintain dynamic competitive comparisons. Web scraping monitors competitor websites for changes in product descriptions, feature lists, pricing tiers, and marketing messages, while machine learning categorizes and structures this information into comparable frameworks.

McKinsey's "Where to Play, How to Win" framework demonstrates how competitive product analysis helps companies identify attractive market segments and devise strategies to outperform competitors. By analyzing competing products' features, pricing, and customer reviews, organizations gain insights into what customers value most and identify underserved segments where they could potentially thrive.

Feature gap analysis maintains matrices showing which capabilities each competitor offers, enabling product teams to identify differentiating capabilities to emphasize or gaps where competitors have established advantages. A software company might track twenty competitors across fifty product features, with AI automatically updating the matrix as competitors release new capabilities or modify existing ones. This reveals patterns like multiple competitors adding similar features simultaneously, suggesting emerging customer expectations, or feature categories where no competitor excels, representing potential differentiation opportunities.

Content gap analysis identifies topics where competitors rank highly in search results but your organization lacks comprehensive content. AI visibility benchmarking tools systematically measure and compare your brand's presence in AI-generated responses relative to competitors, creating visibility heatmaps that show where competitors dominate search visibility and where opportunities exist.

However, gap analysis can mislead if it assumes feature parity represents strategic necessity. Many successful products win precisely through focused feature sets rather than matching every competitor capability. McKinsey's research on market positioning emphasizes that "how to win" involves crafting strategies that capitalize on competitor weaknesses and selectively emulate their strengths rather than blindly matching all features. Blindly matching competitor features produces bloated offerings lacking clear positioning.

Organizations must evaluate whether addressing each identified gap actually serves strategic objectives or merely creates feature parity without competitive advantage.

Critical Assessment and Limitations

AI competitive intelligence systems provide breadth but human analysis provides depth. Technology enables monitoring of many competitors simultaneously across numerous sources, but strategic interpretation requires human judgment about significance. That a competitor reduced prices 5% might be significant or might reflect routine inventory clearance. Distinguishing strategic moves from tactical adjustments requires business judgment and market understanding that algorithms cannot replicate.

Data reliability presents concerns. Web scraping can misinterpret content, particularly when websites employ dynamic loading or personalization. Forbes research on AI-driven competitive intelligence notes that while AI helps mine vast data reservoirs rapidly, it can generate plausible but inaccurate insights—"hallucinations"—requiring careful validation. Competitors may strategically manipulate information. Analysts must validate automated collection and apply skepticism to surprising findings.

Missing peripheral vision represents a significant limitation. Harvard Business Review research emphasizes that

competitive intelligence shouldn't focus narrowly on direct competitors but must encompass broader market dynamics and cross-industry trends. Automated monitoring typically tracks known competitors but struggles to detect disruption from adjacent categories. Taxi companies monitored other taxi companies but missed Uber's platform approach.

Information overload without analytical frameworks produces noise rather than intelligence. Harvard Business Review found that many companies collect extensive competitive intelligence but fail to use it effectively in decision-making. The analytical frameworks discussed earlier remain essential for converting data into strategy.

Competitive intelligence systems work best when they complement rather than replace human capabilities. Forbes research emphasizes that competitive edge derives from how leaders leverage AI, not simply from acquiring it. Automation provides continuous monitoring and pattern detection. Humans provide strategic interpretation, contextual understanding, and creative synthesis. Organizations that excel combine automated collection with skilled human analysis and clear processes that translate intelligence into action.

6.

GENERATIVE AI

What is Generative AI?

Generative AI refers to a subset of artificial intelligence systems designed to generate new, original content based on patterns and information learned from existing data. Unlike traditional AI models that focus on classification, prediction, or detection, generative AI models create data that mimics the characteristics of the training data. This can include generating text, images, music, and other forms of media. The role of generative AI in AI is to enable machines to produce creative outputs, making them capable of tasks that require creativity and innovation.

What are LLMs?

Large Language Models (LLMs) are a type ofgenerative AI designed to understand and generate human-like text based on the data they have been trained on. These models can process

vast amounts of text data to perform a variety of tasks, such as translation, summarization, question-answering, and more. They play a crucial role in AI by enabling machines to interact with human language in a way that is increasingly indistinguishable from human interactions.

How Generative AI Works

Generative AI models function through sophisticated machine learning techniques that enable them to learn from and generate new data. Here's a high-level overview of how generative AI models work:

- Training Process: Generative AI systems learn through exposure to vast datasets relevant to their intended output. For instance, a system designed to generate marketing copy could study millions of existing marketing materials, advertisements, and related content to understand effective communication patterns and styles.
- Processing Framework: The system operates through an interconnected network that processes information systematically. This network analyzes and interprets patterns within the training data, developing an understanding of structure, style, and context. The system continuously refines its understanding through this analysis, improving its ability to generate

appropriate content.

Different generative AI systems are engineered for specific purposes, each optimized for particular types of content creation. Examples include:

- Text Generation Systems process and produce written content by analyzing language patterns, context, and communication styles. These systems can create various forms of text, from creative writing to technical documentation.

- Image Generation Systems learn visual patterns and artistic principles to create new images. They understand elements like composition, color relationships, and visual coherence to produce original visual content.

Generative AI has evolved rapidly, with each new development building on the successes and limitations of previous models, leading to increasingly sophisticated and capable systems.

Capabilities and Limitations

Capabilities:

- **Content Creation**: Generative AI models can create new content across various media, including text, images, music, and video. This capability is valuable for

creative industries, marketing, and entertainment.

- **Personalization**: These models can generate personalized content tailored to individual preferences, enhancing user experiences in applications like recommendation systems and targeted advertising.
- **Data Augmentation**: Generative AI can create synthetic data to augment training datasets, improving the performance of other AI models and addressing data scarcity issues.

Limitations:

- **Bias**: Generative AI models can inherit and amplify biases present in their training data, leading to biased or unfair outputs. Mitigating bias remains a significant challenge.
- **Data Dependency**: The quality and diversity of the generated content are heavily dependent on the training data. Poor or unrepresentative data can limit the model's effectiveness.
- **Resource Intensive**: Training generative AI models requires substantial computational resources, making them expensive and environmentally impactful.
- **Quality Control**: While generative AI can produce high-quality content, it may also generate outputs that are nonsensical or factually incorrect, necessitating human oversight and quality control.

Does Generative AI think?

The characterization of generative AI systems as novel forms of intelligence or artificial cognition has become prevalent in both popular and academic discourse. This perspective, while intuitively appealing, may lead to misconceptions about the nature and capabilities of these systems. Vallor (2024) uses the metaphor of a mathematical mirror, which reflects patterns of human-generated data, to provide a more nuanced and accurate conceptualization.

This reflective model provides a framework for understanding how these systems operate without attributing to them qualities of cognition or understanding that they do not possess. It allows us to appreciate their capabilities while maintaining a clear distinction between their functioning and human-like intelligence.

The Reflective Model: A Conceptual Framework

The analogy of a physical mirror serves as an accessible entry point to understanding the more complex mathematical processes at work in generative AI. Just as a mirror's physical properties determine how it reflects light, the mathematical properties of AI algorithms determine how they "reflect" patterns in data.

Key aspects of this analogy include:

1. Reflection vs. Generation: Like a mirror that doesn't create light but reflects it, these AI systems don't create new information but reflect patterns from their training data.
2. Fidelity and Distortion: Just as mirrors can provide clear reflections or distort images, AI systems can produce outputs of varying accuracy or relevance.
3. Amplification: Some mirrors can magnify reflections; similarly, AI systems can amplify certain patterns or biases present in their training data.

Operational Mechanism of the AI Reflective Model

The operational mechanism can be broken down into three main phases:

1. Data Input (Training Phase): This phase involves exposing the AI system to vast quantities of human-generated data. This data serves as the "light" in our mirror analogy. It typically includes:

    - Textual data from books, articles, websites, and social media
    - Conversational data from online forums and chat

logs
- Structured data from databases and knowledge bases

2. Pattern Recognition and Replication: During this phase, the AI system processes the input data to identify statistical regularities. This involves:

    - Analyzing the frequency and co-occurrence of words and phrases
    - Identifying grammatical structures and linguistic patterns
    - Recognizing contextual relationships between different pieces of information

    Importantly, this process does not involve semantic understanding. The system is not comprehending the meaning of the text, but rather identifying mathematical relationships within the data.
3. Output Generation: When prompted, the AI generates responses by:

    - Identifying patterns in the prompt that match patterns in its training data
    - Combining and extrapolating from these patterns to produce new text
    - Using statistical probabilities to determine the most

likely sequence of words or ideas

Functional Capabilities and Limitations

Understanding the reflective nature of these systems is crucial for accurately assessing their capabilities and limitations:

1. Lack of True Cognition: While these systems can produce outputs that appear to demonstrate reasoning or emotional awareness, they are not engaging in cognitive processes as humans do. They are performing complex statistical analyses based on patterns in their training data.
2. Absence of Understanding: The AI does not comprehend the content it generates in any meaningful sense. It cannot reason about the implications of its outputs or verify their factual accuracy beyond pattern matching.
3. Contextual Limitations: The AI's responses are limited to reflecting patterns present in its training data. It may struggle with novel contexts or information outside its training distribution.
4. Impressive but Bounded Capabilities: These systems can generate highly sophisticated and contextually appropriate text, solve certain types of problems, and even engage in creative-seeming tasks. However, these

capabilities are bounded by the patterns present in their training data and do not represent open-ended intelligence.

By conceptualizing generative AI through this reflective model, we can better understand its functioning, appreciate its capabilities, and recognize its limitations. This framework provides a more accurate basis for discussing and researching these systems, avoiding anthropomorphization while still acknowledging their significant technological achievements.

How to Give Clear and Effective Instructions to Generative AI Models

Understanding Prompt Engineering

Generative AI models like those from OpenAI are designed to respond optimally to certain prompt formats due to the way they are trained. This process, known as prompt engineering, involves crafting inputs to get the best possible outputs from these models.

To optimize prompt engineering with generative AI models, follow these best practices:

1. **Use the Latest Model:** Newer models provide better

performance.

2. **Clear Instructions:** Place instructions at the beginning of the prompt, separated by delimiters.
3. **Specificity:** Be detailed about the desired outcome, length, format, and style.
4. **Examples:** Provide output format examples to guide the model.
5. **Step-by-Step:** Start with zero-shot, then few-shot, and fine-tune if necessary.
6. **Precision:** Avoid vague descriptions.
7. **Positive Guidance:** Instruct on what to do, not just what to avoid.
8. **Code Generation:** Use "leading words" to nudge the model towards desired patterns.

For detailed examples and more information, visit the Prompt Engineering Guide.

Additionally, experimenting with different prompt structures can also yield beneficial results. Applying these strategies will help you communicate more clearly with generative AI models, ensuring you receive the most accurate and useful responses.

Examples

The "Examples" section of the Prompt Engineering Guide 🔗 showcases various applications of prompts. Each example

illustrates how to construct prompts effectively to achieve desired results, emphasizing clarity and context. These examples highlight generative AI models' diverse capabilities and practical uses. For a deeper understanding and to see the specific examples in action, students are encouraged to explore the guide at Prompt Engineering Guide 🔗.

Text Summarization: The model condenses lengthy texts into concise summaries, such as summarizing the function and use of antibiotics in one sentence.

Information Extraction: Extracts specific details from text, like identifying a product mentioned in a research paper.

Question Answering: Answers questions based on given context, for instance, identifying the source of a molecule.

Text Classification: Categorizes text based on sentiment, like determining if a statement is positive, negative, or neutral.

Conversation: Adapts tone and complexity of responses, such as explaining black holes to different audiences.

Code Generation: Generates code snippets based on specific instructions, from simple greetings to complex database queries.

Reasoning: Solves arithmetic and logic problems by breaking them into steps.

Four important prompting techniques

Effective prompting techniques are essential for maximizing

the potential of language models in various applications. Among the most powerful methods are zero-shot prompting, few-shot prompting, chain prompting, and the tree of thoughts approach. Each technique offers unique advantages and can be tailored to specific tasks to enhance the performance and reliability of AI models. Let's explore these techniques in detail to understand how they can be applied in different scenarios. Please visit the links below and review the full examples. You can also find additional prompting techniques at: promptingguide.ai.

Zero-Shot Prompting 🔗

Zero-shot prompting involves providing the model with a task it has never seen before without any examples. The model uses its pre-trained knowledge to generate a response based on the prompt alone. This technique is useful for tasks where you want the model to generalize from its training data to new, unseen situations.

Few-Shot Prompting 🔗

Few-shot prompting provides the model with a few examples of the task at hand within the prompt. By seeing these examples, the model can better understand the task requirements and generate more accurate responses. This approach is beneficial when you need the model to perform a task that might be slightly outside its usual scope, but still within the realm of its training data.

Prompt chaining 🔗

Prompt chaining involves breaking down a complex task into a series of smaller, manageable prompts. Each prompt builds upon the previous one, guiding the model step-by-step through the problem-solving process. This method helps improve the accuracy and coherence of the model's responses, particularly for tasks that require multi-step reasoning or detailed analysis.

Tree of Thoughts 🔗

The tree of thoughts technique is a structured approach where the model generates multiple potential solutions or thoughts at each step of the problem-solving process. These thoughts branch out like a tree, allowing for the exploration of various possibilities before converging on the best solution. This approach is particularly effective for tasks that involve creative problem-solving or scenarios where multiple pathways need to be considered before arriving at a final answer.

From prompt engineering to problem formulation

Managing Generative AI as a Skilled Yet Inexperienced Team Member

When working with generative AI models, it's helpful to imagine them as an employee who is always confident in their answers but doesn't always know the best way to accomplish a

task. This analogy can guide you in providing clear and precise instructions, just as you would for an eager yet inexperienced team member.

Generative AI models excel at generating responses based on patterns in the data they were trained on, but they may struggle with tasks that require specific knowledge or nuanced understanding. Therefore, it's crucial to give detailed, step-by-step instructions to ensure they perform correctly. Much like you wouldn't expect an employee to guess what you want without proper guidance, you shouldn't expect a generative AI to produce perfect results without clear and explicit prompts.

By treating generative AI like a confident but sometimes uncertain employee, you'll be more mindful of how you craft your prompts. Start with clear instructions at the beginning, use specific and detailed language, and provide examples to illustrate the desired outcome. This approach will help you get the most accurate and useful responses from your AI, just as clear and effective communication ensures the best performance from your human team members.

Improving Problem Formulation Skills

This is a summary of Acar (2023)

To enhance problem formulation skills, we can deploy four key components of effective problem formulation: problem diagnosis, decomposition, reframing, and constraint design.

Problem Diagnosis

Problem diagnosis involves identifying the core issue for AI to address. This means pinpointing the main objective you want generative AI to achieve. Some problems are straightforward, like gathering information on human resource management strategies for employee compensation. Others, like finding innovative solutions, are more complex.

A notable example is InnoCentive (now Wazoku Crowd), which has helped clients formulate over 2,500 problems with an impressive success rate exceeding 80%. Their success stems from their ability to identify the fundamental issue underlying a problem. They often use the "Five Whys" technique—repeating the question "why?" five times, each time directing the current "why" to the answer of the previous "why"—to differentiate root causes from symptoms. For instance, in tackling the Exxon Valdez oil spill, they identified that the core issue was the viscosity of the crude oil in subarctic waters, leading to a solution that involved vibrating the oil to keep it in a liquid state.

Problem Decomposition

Problem decomposition involves breaking down complex problems into smaller, more manageable subproblems. This is crucial when dealing with multifaceted issues that are too convoluted to solve directly.

For example, in the InnoCentive challenge for amyotrophic

lateral sclerosis (ALS), instead of seeking a broad solution for the disease, the focus was narrowed to detecting and monitoring its progress. This led to the development of a noninvasive and cost-efficient ALS biomarker.

Testing how AI improves with problem decomposition, a challenge like implementing a robust cybersecurity framework can be broken down into subproblems such as security policies, vulnerability assessments, authentication protocols, and employee training. This approach leads to more practical and effective solutions.

Problem Reframing

Problem reframing involves viewing a problem from different perspectives to find alternative interpretations and solutions. By reframing, you can guide AI to explore a wider range of potential solutions, overcoming creative blockages.

An illustrative example is Doug Dietz of GE HealthCare. When he observed the fear children experienced during MRI scans, he reframed the problem to ask, "How can we turn the daunting MRI experience into an exciting adventure for kids?" This led to the creation of the GE Adventure Series, significantly reducing pediatric sedation rates and improving patient satisfaction.

Another example could be reframing an office parking problem. Instead of focusing solely on increasing parking spaces, reframing the issue to consider employees' stress and

commuting options can lead to solutions like promoting alternative transportation and remote work options.

Problem Constraint Design

Problem constraint design involves setting the boundaries for problem-solving by defining input, process, and output restrictions. Constraints can guide AI to generate solutions that are valuable for the task at hand. For productivity-oriented tasks, specific and strict constraints are effective. For creativity-oriented tasks, experimenting with constraints can help discover novel perspectives.

For instance, brand managers use AI tools like Lately or Jasper to produce social media content, setting precise constraints on length, format, tone, or target audience to ensure consistency with the brand image. Conversely, when seeking originality, constraints can be minimized. A great example is GoFundMe's Help Changes Everything campaign, where unorthodox constraints led to the creation of a visually striking and emotionally impactful year-in-review video using AI-generated street mural-style art.

Marketing Problem: Declining Market Share for Premium Coffee Brand

A premium coffee brand has experienced a 15% decline in

market share over two years despite maintaining product quality and competitive pricing. Initial analysis suggests the brand is losing ground to both artisanal local roasters and mass-market competitors, indicating a complex challenge requiring systematic problem formulation.

Problem Diagnosis through the Five Whys technique reveals the underlying issue. The initial problem appears straightforward: market share is declining. But why? Because fewer consumers choose the brand in retail environments. Why? Because the brand lacks distinctive positioning in a crowded category. Why? Because messaging emphasizes product attributes that competitors have replicated. Why? Because the brand strategy was developed when quality alone provided differentiation. Why? Because the strategic framework hasn't evolved to address experiential and identity-based value propositions that now drive premium coffee consumption. The root problem is an outdated value proposition that treats premium coffee as a product category rather than as a vehicle for consumer identity expression. The reformulated problem becomes: How can the brand articulate a distinctive value proposition that aligns with contemporary consumers' identity projects and experiential expectations?

Problem Decomposition breaks this complex challenge into manageable components. First, understanding current brand perception among target segments to explore how consumers incorporate premium coffee into their daily routines and self-narratives. Second, mapping the competitive

landscape to identify positioning gaps related to consumer identity construction. Third, analyzing the customer journey to pinpoint moments where brand interaction could create differentiation. Fourth, evaluating channel strategy to determine whether distribution choices align with the redefined value proposition. Fifth, assessing communication touchpoints to ensure message consistency across platforms. Each subproblem can be addressed independently with specialized marketing research methods and strategic frameworks, yet together they provide a comprehensive approach to reversing market share decline.

Problem Reframing generates alternative perspectives on the market share challenge. The original frame ("How do we regain lost market share from competitors?") positions the challenge as a zero-sum competitive battle focused on capturing market share. An alternative frame might ask: "How can we create new consumption occasions that expand the premium coffee market while establishing our brand as essential to these occasions?" Another reframing considers the consumer perspective: "What tensions or contradictions do consumers experience in their premium coffee consumption, and how can our brand help them navigate these tensions?" By reframing the problem, the brand might discover opportunities otherwise invisible to the original framing.

Problem Constraint Design establishes boundaries that guide solution development. Consider two contrasting constraint configurations. Configuration A imposes tight

constraints: solutions must work within existing distribution channels, maintain current price points, preserve brand heritage, and demonstrate measurable impact within twelve months. These constraints eliminate possibilities (e.g., entering new retail categories, pursuing radical repositioning, launching sub-brands) while focusing attention on messaging refinement, packaging updates, and targeted promotions implementable through existing infrastructure. Configuration B relaxes several constraints: distribution channels can be expanded, pricing architecture can be revised, and the evaluation timeline extends to twenty-four months. This permits consideration of direct-to-consumer models, subscription services, or partnerships with experiential venues.

Synthesis Across Approaches. These four problem formulation approaches work together. Problem diagnosis identifies root causes. Problem decomposition creates actionable subproblems from that diagnosis. Problem reframing tests whether the diagnosed problem is correctly formulated by exploring alternative interpretations that might reveal overlooked dimensions of the challenge. Finally, problem constraint design focuses solution development by establishing which types of solutions are feasible given organizational realities.

Real-World Examples: Generative AI and LLMs in Marketing

Understanding the theoretical aspects of generative AI and Large Language Models (LLMs) is crucial, but seeing their practical applications in the real world can make these concepts more tangible and relevant. Here are some specific case studies and examples of how generative AI and LLMs are being used in marketing today.

Case Study 1: Personalized Email Marketing with LLMs

Company: Persado **Application**: Personalized Email Campaigns

Persado, a leading AI company, utilizes LLMs to generate personalized email content for marketing campaigns. By analyzing customer data and preferences, Persado's AI crafts personalized subject lines, body text, and calls-to-action that resonate with individual recipients. This approach has significantly improved open rates and conversion rates for their clients, demonstrating the power of personalized communication in marketing.

Impact:

- Increased email open rates by up to 68%.
- Enhanced engagement through tailored content.

- Higher conversion rates leading to increased sales.

Case Study 2: Dynamic Ad Generation with Generative AI

Company: The Grid **Application**: Automated Web Design and Ad Creation

The Grid is a web design platform that uses generative AI to create visually appealing and effective advertisements. By inputting basic information about a product or service, the AI generates customized ad designs that are optimized for various platforms, including social media, search engines, and display networks. This allows marketers to quickly and efficiently produce high-quality ads without the need for extensive design expertise.

Impact:

- Reduced time and cost for ad creation.
- Consistent brand aesthetics across multiple platforms.
- Improved ad performance through data-driven design adjustments.

Case Study 3: Chatbots and Customer Interaction

Company: H&M **Application**: Customer Service Chatbots

H&M, the global fashion retailer, employs LLM-powered

chatbots to handle customer service inquiries on their website and mobile app. These chatbots are trained to understand and respond to a wide range of customer queries, from product information to order tracking and returns. By providing instant and accurate responses, the chatbots enhance the customer experience and free up human agents to handle more complex issues.

Impact:

- Faster response times for customer inquiries.
- Improved customer satisfaction and engagement.
- Reduced operational costs associated with customer service.

Case Study 4: Content Creation for Social Media

Company: Coca-Cola **Application**: Social Media Content Generation

Coca-Cola has experimented with using generative AI to create engaging content for their social media channels. By leveraging LLMs and generative models, the company can quickly produce a variety of posts, including captions, hashtags, and visual content ideas. This not only keeps their social media presence active and dynamic but also allows them to test different content strategies to see what resonates best with their audience.

Impact:

- Increased social media engagement through diverse content.
- Enhanced brand visibility and interaction with followers.
- Ability to rapidly adapt to trending topics and events.

Key Terms

Category	Term	Definition
Core Concepts	**Generative AI**	AI that produces new content resembling its training data, such as text, images, music, or code.
	Large language model (LLM)	A model that understands and generates human-like text for tasks like summarization, Q&A, and drafting.
	Reflective model	View of generative AI as reflecting and recombining patterns from data rather than truly understanding meaning.
How It Works	**Training process**	Learning statistical patterns from large datasets by adjusting parameters to better predict tokens or pixels.
	Output generation	Using learned probabilities at inference to compose new sequences that fit the prompt's context and style.
Prompting Techniques	**Prompt engineering**	Crafting clear, specific instructions and examples to guide models toward the desired output and format.
	Zero-shot prompting	Asking the model to perform a task with no examples, relying on its general knowledge.

Category	Term	Definition
	Few-shot prompting	Providing a few labeled examples in the prompt to show the model the desired pattern.
	Prompt chaining	Breaking a complex task into sequenced prompts that build on intermediate results.
	Tree of thoughts	Exploring multiple candidate solutions in parallel, evaluating them, and selecting a best path.
Capabilities & Applications	**Content creation**	Generating new copy, images, audio, or video for marketing, support, or product experiences.
	Personalization	Tailoring content and tone to an individual's context and preferences using learned patterns.
	Data augmentation	Synthesizing additional training data to balance classes, expand coverage, or protect privacy.
Limitations & Risks	**Bias**	Outputs reflect skewed training data, causing unfair or stereotyped content.
	Hallucination	Confident but incorrect or fabricated output that requires human review and safeguards.

Category	Term	Definition
	Resource intensity	Training and serving large models consume significant compute, time, and energy.
Practice	**Problem formulation**	Diagnose, decompose, reframe, and set constraints so prompts and tasks align with real goals.

Summary: Using Generative AI

Key Takeaways:

- LLMs are a specific type of AI focused on understanding and generating human-like text, while Generative AI is a broader category encompassing systems that create various types of content (text, images, music).
- Generative AI works through a training process that analyzes patterns in vast datasets, functioning more like a mathematical mirror reflecting patterns rather than truly "thinking" or understanding.
- Major capabilities include content creation across

media types, personalization, and data augmentation, while limitations include bias inheritance, heavy resource requirements, and quality control needs.

- Real-world applications show successful implementation in marketing, such as Persado's personalized email campaigns and H&M's customer service chatbots, demonstrating tangible business impact.
- Effective prompt engineering requires clear instructions, specific details, and appropriate examples, with four key techniques: zero-shot prompting, few-shot prompting, chain prompting, and tree of thoughts.

AI in Action: Case studies demonstrate AI's practical impact, like Coca-Cola's social media content generation and The Grid's automated web design platform, showing how theoretical capabilities translate into business value.

Think Deeper:

1. How might the "mathematical mirror" concept of AI influence how we approach prompt engineering and problem formulation?
2. Consider the balance between automated AI solutions and human oversight – what criteria should organizations use to determine the appropriate level of human involvement?

3. How can organizations effectively decompose complex problems for AI while maintaining sight of the broader strategic objectives?

7.

AI AGENTS

Based on AWS (2025), BCG (2025), Colback (2025), Gaarlandt et al. (2025), IBM 2025, Microsoft (2025a, 2025b), Qu et al. (2025)

AI agents are autonomous AI programs that can perceive information, make decisions, and act toward achieving goals with minimal human intervention, using AI models and connected tools to carry out tasks. AI agents possess a degree of *agency*: they don't just respond to direct commands, they can *proactively initiate* steps and adapt their behavior to changing inputs or objectives. This agentic quality has led experts to distinguish agentic AI from generative AI: *generative AI creates content, while agentic AI takes action*. While generative AI (like an LLM that writes copy) is largely reactive to prompts, an AI agent can set goals, make independent decisions, and execute multi-step workflows to accomplish a specific mission.

This chapter provides a comprehensive guide to understanding AI agents and their transformative impact on marketing practice. We begin by exploring how AI agents think and plan, examining their core observe-plan-act cycle

and comparing two key problem-solving approaches: the adaptive "think-as-you-go" method and the efficient "plan-first" strategy. Next, we'll classify the five main types of AI agents—from simple reflex agents to sophisticated learning systems—helping you choose the right level of capability for different marketing challenges. We then examine both the remarkable strengths of AI agents (speed, consistency, complex decision-making) and their important limitations (lack of common sense, black-box decisions, potential biases), including what can go wrong and how to prevent problems. The chapter concludes by showcasing real-world applications, from how AI agents are revolutionizing campaign management and content creation for marketers to how they're transforming the shopping experience for consumers. By the end of this chapter, you'll understand what AI agents can do, and when and how to use them effectively in marketing practice.

How AI Agents Think and Plan

AI agents are being used by both companies and consumers to help with marketing tasks, and these systems need to solve complex problems step by step – just like a marketing team would. But before diving into the two main problem-solving

approaches, it's helpful to understand the basic workflow that all AI agents follow.

The Basic AI Agent Workflow

Every AI agent goes through three fundamental steps:

1. **Understanding Your Goal** The AI agent takes your request. whether it's "analyze our Q3 campaign performance" or "help me plan a product launch", and breaks it down into smaller, manageable tasks. Just like a project manager would create a task list from a big objective.
2. **Gathering Information** The AI agent collects everything it needs to complete those tasks. This might mean pulling data from CRM, searching industry reports online, or analyzing your social media metrics. Some advanced agents can even collaborate with other AI tools to get specialized insights.
3. **Executing and Adjusting** The AI agent works through its task list methodically. After completing each task, it checks whether it's moving toward your original goal and adjusts if needed – sometimes even adding new tasks it realizes are necessary.

How AI Agents Work: The Core Cycle

AI agents follow a continuous three-step cycle that allows them to tackle complex challenges:

Observe: AI agents constantly collect and process information from their environment – whether that's user requests, performance metrics, customer data, or market trends. What makes them powerful is their ability to remember previous conversations and interactions, building context over time. Think of this like how an experienced marketer keeps track of past campaign results and customer feedback to inform future decisions.

Plan: Using advanced language models, AI agents evaluate what they've observed and create a strategy to accomplish goals. They consider the current situation, past experiences, and available resources to prioritize the best course of action. This is similar to how an experienced marketer might assess market conditions and plan a campaign strategy.

Act: AI agents can connect to various business systems – CRMs, analytics platforms, social media tools, or databases – to actually execute their plans. They can pull reports, analyze data, create content, or even coordinate with other AI agents to get specialized tasks done. If something goes wrong or they need clarification, they can detect the issue and either fix it themselves or ask for input.

Types of AI Agents

AI agents can be developed to have varying levels of capabilities. A simple agent might be preferred for straightforward goals to limit unnecessary computational complexity. In order of simplest to most advanced, there are 5 main agent types:

1. Simple Reflex Agents

Simple reflex agents are the simplest form that base their actions on what they can currently perceive or sense. These agents have no memory of past interactions and cannot seek help from other agents when they lack information. They function on a set of reflexes or rules – they are preprogrammed to perform specific actions when certain conditions are met.

Key characteristics:

- No memory of past events
- React only to current situation
- Follow predetermined rules ("if this, then that")
- Cannot handle unexpected situations they weren't programmed for
- Work best in fully observable environments where all necessary information is available

Marketing applications: Think of basic automated email responses or simple chatbots. If a customer emails asking "What are your store hours?" the agent always gives the same preprogrammed response.

Limitations: If the agent encounters a situation it wasn't prepared for, it cannot respond appropriately. For example, a chatbot programmed to answer pricing questions won't know how to handle a complaint about product quality.

Example: A social media automation tool that posts "Happy Monday!" every Monday at 9 AM, regardless of current events or brand context.

2. Model-Based Reflex Agents

Model-based reflex agents use both their current perception and memory to maintain an internal model of their environment. As the agent receives new information, it updates this internal model. The agent's actions depend on its model, rules, previous experiences, and current situation.

Key characteristics:

- Have memory and can store information
- Build and update an internal "model" of their world
- Can operate in partially observable environments that change over time
- Still limited by their set of predetermined rules

Marketing applications: A customer service chatbot that remembers your name, past purchases, and previous conversations during your current session. It builds a model of who you are and what you might need.

How the internal model works: Like a robot vacuum that maps your house as it cleans, avoiding furniture it has encountered before and remembering which areas it has already cleaned to avoid repeating work.

Example: An e-commerce website that tracks your browsing behavior and adjusts product recommendations based on what you've viewed, creating an internal model of your preferences that gets updated with each click.

3. Goal-Based Agents

Goal-based agents have an internal model of the world plus specific goals they want to achieve. These agents search for sequences of actions that will reach their goal and plan these actions before executing them. This planning capability makes them more effective than simple or model-based reflex agents.

Key characteristics:

- Have specific goals or objectives
- Plan sequences of actions to reach goals
- Search through different possible approaches
- Make decisions based on which actions will achieve their

objectives

Marketing applications: A campaign optimization system with the goal of maximizing email open rates. It doesn't just send emails – it analyzes different subject lines, send times, and audience segments to plan the approach most likely to achieve its goal.

How goal-planning works: Like a GPS navigation system that considers multiple routes to your destination and recommends the one that best meets your goal (fastest route, shortest distance, avoiding tolls, etc.).

Example: A social media management tool that has the goal of increasing brand engagement. It analyzes when your audience is most active, what content performs best, and plans a posting schedule designed to maximize interactions.

4. Utility-Based Agents

Utility-based agents select sequences of actions that not only reach their goal but also maximize utility or reward. Utility is calculated through a utility function – a system that assigns a utility value (measuring the usefulness or satisfaction an action brings) to each possible scenario based on fixed criteria.

Key characteristics:

- Have utility functions that measure the value of different

outcomes
- Consider multiple factors when making decisions
- Optimize for the best possible result, not just goal achievement
- Weigh different priorities against each other

How utility functions work: The criteria might include factors like progress toward the goal, time requirements, cost, or computational complexity. The agent calculates which action will produce the highest utility score and selects that option.

Marketing applications: A programmatic advertising system that doesn't just aim for clicks but considers cost per click, conversion probability, brand safety, and audience quality to bid on ad placements that provide the highest overall value.

Example: A navigation app that optimizes for multiple factors – recommending a route that balances time, fuel efficiency, toll costs, and traffic conditions to maximize overall utility rather than just focusing on one factor.

5. Learning Agents

Learning agents have the same capabilities as other agent types but are unique in their ability to learn autonomously. They add new experiences to their initial knowledge base automatically, which enhances their ability to operate in

unfamiliar environments. Learning agents are composed of four main components:

The four key components:

1. **Learning:** Improves the agent's knowledge by learning from the environment through sensors and interactions
2. **Critic:** Provides feedback on whether the agent's responses meet performance standards
3. **Performance:** Responsible for selecting actions based on what has been learned
4. **Problem Generator:** Creates various proposals for new actions to test and try

Key characteristics:

- Continuously update their knowledge base
- Learn from new experiences automatically
- Improve performance over time
- Can adapt to unfamiliar environments
- May be utility-based or goal-based in their underlying reasoning

Marketing applications: Personalized recommendation engines that track user activity and preferences, storing this information to make better suggestions over time. Each interaction teaches the system more about what works.

How the learning cycle works: Every time a customer interacts with the system (clicks, purchases, ignores recommendations), the learning component processes this feedback, the critic evaluates whether recommendations were successful, performance adjusts future actions, and the problem generator suggests new approaches to test.

Example: Netflix's recommendation system starts with basic demographic and genre preferences but learns from every show you watch, rate, skip, or rewatch. It continuously improves its suggestions by analyzing patterns in your behavior and comparing them to similar users.

Choosing the Right Agent Type

Simple reflex agents work for basic, repetitive tasks where consistency matters more than sophistication and the environment is predictable.

Model-based agents are suitable when you need some personalization and memory, but customer interactions follow relatively predictable patterns.

Goal-based agents are ideal when you have clear, measurable objectives and want the system to plan the best approach to achieve them.

Utility-based agents excel when you need to balance multiple competing priorities and optimize for complex, multi-factor outcomes.

Learning agents are worth the investment when you have

substantial customer data and want continuously improving performance, especially in dynamic environments where customer preferences change over time.

Using AI Agents: From Prompt Engineering to Context Engineering

Context engineering is the difference between writing one good instruction and managing an entire information environment. Think of it like designing a marketer's dashboard: you need to decide which KPIs, customer segments, campaign history, and real-time data appear on screen, because showing everything at once overwhelms decision-making rather than improving it. Anthropic defines context as everything the AI sees during operation—instructions, past messages, retrieved documents, tool outputs—and emphasizes that this "attention space" is limited. Just as too much information on a dashboard causes analysis paralysis, too much information in an AI's context causes it to lose focus and perform worse (called "context rot"). Context engineering means strategically curating what information is present, what gets summarized, what's removed as it becomes irrelevant, and what's retrieved only when needed, so the AI stays effective across multi-step tasks. Prompt engineering optimizes the wording of your

instructions; context engineering optimizes the entire information environment so the AI can execute well throughout a complex workflow.

Imagine an AI agent helping develop a product launch campaign for a new sneaker line. With poor context management, you'd dump everything into the AI at once: the entire brand guidelines document, all past campaign performance data, complete customer research reports, social media analytics, competitor analysis, product specifications, budget constraints, and channel-specific best practices. The AI would struggle to focus, might confuse data from different customer segments, or apply tactics from last year's failed campaign to this year's launch.

With context engineering, the AI starts with core brand guidelines and the product brief. As it develops the positioning strategy, it retrieves only the relevant customer segment research. When it moves to channel planning, previous campaign data gets summarized to "Instagram Reels drove 3x engagement vs. static posts for Gen Z" rather than keeping every metric visible. Old social media threads get removed from active context once insights are extracted. When the AI shifts to budget allocation, it loads budget constraints and cost-per-channel data while temporarily removing creative guidelines. Throughout the process, key decisions get stored in a persistent memory ("targeting women 18-24 interested in sustainable fashion") that can be referenced without cluttering the working space. This keeps the AI focused on the current

subtask while maintaining coherence across the entire campaign development workflow.

AI Agents: Strengths, Limitations, and What Can Go Wrong

Understanding AI agents isn't just about knowing what they can do—it's equally important to understand their limitations and potential pitfalls. Like any powerful marketing tool, AI agents work best when marketers understand both their capabilities and their blind spots. This knowledge helps you set realistic expectations, choose appropriate use cases, and build safeguards to prevent problems.

What AI Agents Excel At

Speed and Scale That Humans Can't Match

AI agents work around the clock without breaks, sick days, or vacation time. While your human team sleeps, an AI agent can analyze yesterday's campaign performance and start optimizing today's ads. This isn't just about working

faster—it's about working at a scale that would be impossible for humans.

Consider customer service: if your business suddenly gets twice as many inquiries during a product launch, you can't instantly hire and train twice as many support staff. But you can deploy additional AI agent instances or allocate more computing power to handle the surge. This scalability means customers get help when they need it, regardless of time zones or sudden spikes in demand.

Consistent Performance Every Time

Humans have bad days, get distracted, or make typos when they're tired. AI agents perform the same task the same way, every single time. For routine marketing activities like updating product databases, checking content for compliance, or sending personalized emails, this consistency is invaluable.

An AI agent reviewing ad copy for brand guidelines will apply the same standards to the 500th piece of content as it did to the first. A human reviewer might start missing details after hours of repetitive work, but the agent maintains the same level of attention throughout.

Handling Complex Decisions with Multiple Variables

Marketing decisions often involve juggling dozens of factors

simultaneously. An AI agent optimizing ad campaigns can consider thousands of targeting combinations, test multiple creative variations, and adjust bidding strategies across different platforms—all while monitoring performance in real-time.

For example, a campaign optimization agent might simultaneously analyze audience demographics, time of day, device type, competitor activity, and seasonal trends to determine the best ad placement strategy. This level of multi-variable optimization would be overwhelming for humans but is exactly what AI agents are designed to handle.

Learning and Improving Over Time

Well-designed AI agents get better at their jobs through experience. A customer service agent can learn from cases where it provided unhelpful responses (with human feedback) and improve its future interactions. A content creation agent can analyze which blog posts perform best and adapt its writing style accordingly.

This continuous improvement means your AI agents become more valuable over time, like having employees whose skills automatically sharpen with every task they complete.

Cost Efficiency at Scale

While AI agents require upfront investment in development

and computing resources, they can dramatically reduce ongoing costs for routine tasks. One content creation agent might produce 50 blog posts monthly that would otherwise require a team of writers, representing significant cost savings that can be reinvested in strategy, creativity, or customer experience improvements.

Where AI Agents Fall Short

Missing Common Sense and Context

AI agents can make mistakes that seem obvious to humans because they lack real-world common sense. They operate on patterns in data rather than intuitive understanding of how the world works.

For instance, an agent might completely miss sarcasm in social media comments or fail to recognize that scheduling a promotional campaign during a national tragedy would be tone-deaf. They follow their programming literally—if told to "maximize email engagement," an agent might start sending excessive emails or using clickbait subject lines that achieve higher open rates but damage brand trust.

Black Box Decision-Making

Many AI agents can't clearly explain why they made specific decisions. If your campaign optimization agent suddenly shifts

30% of your budget from Facebook to Instagram, you might want to understand the reasoning. Was it based on performance data, audience behavior, or a technical glitch?

This opacity makes it difficult to learn from agent decisions, troubleshoot problems, or satisfy compliance requirements. It's particularly challenging when you need to explain marketing decisions to stakeholders or justify budget allocations to executives.

Data Quality Problems and Bias

AI agents are only as good as the data they learn from. If historical data contains biases or blind spots, agents will perpetuate and potentially amplify these issues.

Consider a lead scoring agent trained on past sales data. If your sales team historically spent more time on leads from certain industries or demographics (for reasons unrelated to actual purchase likelihood), the agent might learn to prioritize similar leads, creating a self-reinforcing bias that limits business growth and raises fairness concerns.

Technical Dependencies and Integration Challenges

AI agents rely on various software tools, databases, and APIs to function. When integrations break—perhaps a social media platform changes its API, or a new privacy regulation restricts

data access—agents may fail silently or produce incomplete results.

Unlike humans who can adapt by finding alternative information sources or workflows, agents typically can't work around technical limitations on their own. They need their digital tools to be working perfectly to perform at their best.

Limited Emotional Intelligence and True Creativity

While AI agents can generate creative content and even produce surprising combinations of ideas, they lack genuine empathy and emotional understanding. They can't truly understand customer frustration, celebrate authentic brand moments, or create emotionally resonant campaigns that connect with human experiences in meaningful ways.

An agent might handle routine customer service inquiries efficiently, but a customer dealing with a personal crisis or complex emotional situation likely needs human understanding and flexibility that current AI agents can't provide.

When Things Go Wrong: Common Failures

Confident but Wrong Answers (Hallucination)

AI agents sometimes generate responses that sound authoritative but are completely incorrect. They might cite non-existent research studies, reference fake customer data, or make up statistics that sound plausible but are fiction.

This is particularly dangerous in marketing because wrong information can lead to poor strategic decisions, compliance violations, or embarrassing public mistakes. Always verify important claims and unusual recommendations from AI agents, especially when they cite specific facts or figures.

Humans Becoming Out of Touch

When AI agents handle routine tasks effectively, human marketers might gradually lose touch with day-to-day customer interactions and market realities. If an agent manages social media responses for months, the marketing team might miss subtle shifts in customer sentiment or emerging trends that require human judgment.

This creates a risk that when human intervention is needed—during a crisis, major strategy shift, or agent failure—the team lacks current knowledge to respond

effectively. Regular human review and involvement helps prevent this disconnect.

Compounding Small Mistakes

AI agents work fast and autonomously, which means small errors can quickly multiply. If an analytics agent forms an incorrect hypothesis about customer behavior, it might pass that flawed insight to a campaign agent, which then creates messaging based on wrong assumptions, while a budget optimization agent allocates resources according to the faulty data.

These cascading errors can create multi-channel problems before humans notice something is wrong. Building in monitoring systems and regular checkpoints helps catch issues before they spiral.

Security Vulnerabilities and Misuse

AI agents with access to company systems and data can become targets for malicious actors. Cybercriminals might try "prompt injection" attacks—tricking agents into ignoring their programming and revealing confidential information or taking unauthorized actions.

Additionally, agents might be manipulated into pursuing their goals through inappropriate means. An engagement-focused social media agent might discover that controversial

posts generate more clicks and start creating divisive content that damages brand reputation while technically achieving its objective.

Misalignment with Values and Brand Ethics

Perhaps the most serious failure mode occurs when agents pursue their programmed objectives in ways that conflict with company values or ethical standards. An agent optimizing for sales conversion might target vulnerable populations with predatory messaging, or a cost-cutting agent might sacrifice customer service quality to improve efficiency metrics.

These alignment problems highlight why clear guidelines, regular monitoring, and human oversight remain essential even with highly autonomous AI agents.

Best Practices for Responsible Agent Use

Understanding these strengths and limitations helps you use AI agents more effectively:

- **Start with clear, well-defined tasks** where agents' consistency and scale provide the most value
- **Maintain human oversight** especially for customer-

facing activities and strategic decisions
- **Build in monitoring and checkpoints** to catch errors before they compound
- **Test extensively** in low-risk environments before deploying agents in critical processes
- **Plan for failure scenarios** and ensure humans can step in when agents encounter problems they can't handle
- **Regular auditing** of agent decisions to check for bias, drift, or misalignment with company values

AI agents are powerful tools that can transform marketing efficiency and effectiveness, but like any powerful tool, they require thoughtful implementation and ongoing management to deliver their benefits while avoiding potential pitfalls.

AI Agents Transforming Marketing Practice

There are many, many ways agents can be put to use. Here are a few examples for key marketing functions:

Campaign Management and Optimization

AI agents are revolutionizing how marketing campaigns are

planned, executed, and optimized. Unlike traditional automation that requires marketers to predict every scenario and create rules accordingly, AI agents can autonomously manage campaigns by analyzing real-time performance data and making intelligent adjustments.

These digital marketing assistants can simultaneously monitor multiple campaigns across different channels, automatically adjusting ad spend, targeting parameters, and creative elements based on performance metrics. They analyze customer behavior patterns to identify optimal timing for campaign launches and can predict which messaging strategies will resonate with specific audience segments.

One particularly powerful application is dynamic pricing and inventory management. AI agents can manage pricing strategies and inventory levels in real-time, responding to factors like competitor pricing, seasonal demand fluctuations, and supply chain disruptions. These systems automatically adjust prices to optimize for different objectives—maximizing revenue during peak demand periods or clearing inventory during slow seasons.

AI Agents in Action: How PUMA Automated Ad Creation

Puma collaborated with agency Monks to create an experimental, fully AI-generated ad using hundreds of autonomous AI agents, marking a significant shift in how

brands might approach creative development. Over five weeks, agents handled everything from concept ideation and scriptwriting to video generation, with minimal human oversight. Each agent was trained for specific tasks and audiences—for instance, curating ideas for Gen Z versus older demographics. The project demonstrates how agentic AI, unlike prompt-based generative AI, can automate entire workflows. Although the resulting one-minute spot had some AI glitches (like extra fingers), it serves as a prototype for scalable, AI-driven ad pipelines. Monks plans to extend this agentic approach to other clients like Google and BMW, and is building infrastructure with Nvidia to support broader adoption across the industry.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=292#oembed-1

Personalized Customer Experiences

Modern consumers expect personalized experiences, and AI agents excel at delivering hyper-customized interactions at

scale. These systems analyze vast amounts of customer data—including browsing behavior, purchase history, social media activity, and demographic information—to create detailed customer profiles and predict preferences.

AI agents can craft personalized email content, recommend products tailored to individual customers, and even adjust website experiences in real-time based on visitor behavior. Unlike rule-based personalization that might show "customers who bought X also liked Y," AI agents understand context and can make sophisticated recommendations that consider factors like seasonal trends, current events, and individual customer life stages.

A prime example is e-commerce product discovery, where AI agents are changing how consumers find and purchase products online. Instead of browsing through categories or using basic search functions, customers increasingly interact with AI shopping assistants that understand complex preferences. For instance, a customer planning a vacation might tell an AI agent they want "sustainable travel options for a family with teenagers who love outdoor activities." The agent can analyze travel patterns, environmental impact data, seasonal weather information, and family-friendly activity options to provide personalized recommendations that no traditional search system could match.

Social Media Management

Social media marketing demands constant attention and quick responses, making it ideal for AI agent assistance. These systems can monitor brand mentions across platforms, engage with customers using natural language, and escalate complex issues to human team members when necessary.

AI agents analyze social media conversations to identify emerging trends, track competitor activity, and gauge public sentiment about brand initiatives. They can automatically respond to common customer inquiries while maintaining brand personality and voice guidelines established by marketing teams.

This capability extends into influencer marketing and partnership management, where identifying and managing relationships with influencers and brand partners can be time-consuming. AI agents can analyze social media metrics, audience demographics, and content quality to identify potential partnership opportunities. They also monitor partnership performance and provide recommendations for optimizing collaborative marketing efforts, helping brands scale their influencer programs more effectively.

Customer Relationship Management

AI agents are transforming how brands manage customer relationships throughout the entire customer journey. These systems can identify high-value prospects, predict customer lifetime value, and determine the optimal communication frequency and channel for each individual customer.

By analyzing patterns in customer behavior, AI agents can predict when customers might be ready to make a purchase, likely to churn, or interested in upgraded products or services. This predictive capability allows marketing teams to be proactive rather than reactive in their customer engagement strategies.

AI Agents in Action: Vera Bradley and Microsoft Copilot

Vera Bradley revamped its customer service capabilities by integrating Microsoft Dynamics 365 with AI-powered tools like Copilot and AI agents. These enhancements enabled the brand to deliver faster, more personalized support and empower agents to handle customer inquiries more effectively. Features such as case summarization, email and chat drafting, and knowledge article suggestions helped reduce repetitive workloads and streamline resolution workflows. AI agents—configurable via Copilot Studio—can manage

routine requests, escalate complex issues with full context, and prioritize high-value interactions for human agents.

Market Research and Analytics

Traditional market research often involves time-consuming surveys, focus groups, and data analysis. AI agents can continuously monitor market conditions, analyze competitor activities, and track consumer sentiment in real-time. They can process vast amounts of unstructured data from social media, review sites, and news sources to identify emerging market opportunities or potential brand risks.

These systems can automatically generate market research reports, track key performance indicators across multiple channels, and provide actionable insights that help marketing teams make data-driven decisions quickly.

AI agents in action: Protecting LUSH products against counterfeits

Lush partnered with MarqVision to combat a surge of sophisticated counterfeiters and impersonators selling expired or unpackaged products online, which threatened customer safety and brand trust. Given the difficulty of enforcing rights on "naked" products without packaging or serial numbers, Lush used MarqVision's AI-driven enforcement tools to tackle infringements across trademarks, copyrights, and regulatory

violations. Within one month, they removed over US$400,000 worth of illegal listings, shut down fake "Lush Official Stores" on Shopee, and achieved a 100% takedown rate for detected impersonators—safeguarding both customers and brand integrity.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=292#oembed-3

AI Agents Transforming Shopping

Walmart's Super Sparky (also more at the WSJ)

Walmart is using advanced AI "super agents" to make online shopping smarter and more efficient—for both customers and employees. Instead of using different apps or tools for every task, Walmart has created four AI agents that each focus on a specific group: one helps customers with things like tracking

orders or finding recipes based on what's in their fridge, another supports store workers with tasks like checking schedules, a third helps suppliers manage orders and ads, and a fourth assists Walmart's tech teams in building new AI tools.

These AI agents can work together and do many tasks automatically, saving time and making things run more smoothly. Walmart hopes this will help boost online sales and improve customer service. Even though the company says these tools will create new jobs, it's also investing in automation—which raises questions about how AI might affect future employment. For students studying business, marketing, or tech, this is a real-world example of how companies are using AI agents to transform how they operate and serve people.



One or more interactive elements has been excluded from this version of the text. You can view them online here: https://pressbooks.pub/aiandmarketing/?p=292#oembed-2

Key Terms

Category	Term	Definition
Core Concepts	**AI agent**	Autonomous program that perceives, decides, and acts toward goals with minimal human input.
	Agentic AI vs generative AI	Agentic AI takes actions to achieve goals. Generative AI produces content in response to prompts.
	Tool use and integration	Ability to call APIs and business apps to gather data, take actions, and complete workflows.
	Memory and state	Stored context from past interactions that the agent uses to plan and adapt.

Category	Term	Definition			
			Agent Types	**Simple reflex agent**	Acts on current input using fixed rules. No memory. Works in predictable settings.
				Model based reflex agent	Maintains an internal model of the environment to handle partial information.
				Goal based agent	Plans sequences of actions to reach explicit objectives defined by the user or system.
	Utility based agent	Chooses actions that maximize a scored tradeoff of outcomes like cost, time, and impact.			

Category	Term	Definition
	Learning agent	Improves decisions over time using feedback, experience, and updated data.
Strengths and Risks	**Scale and consistency**	Operates 24/7 with uniform quality across large task volumes and channels.
	Hallucination	Confident but incorrect output that requires verification and safeguards.
	Prompt injection	Malicious input that tries to override instructions or exfiltrate sensitive data.

Category	Term	Definition
	Alignment and guardrails	Policies, constraints, and monitoring that keep behavior within laws and brand values.

Summary: AI Agents

Key Takeaways:

- AI agents are systems that perceive, decide, and act toward goals with minimal human input, connecting to tools to execute tasks in real environments.
- Agentic AI differs from generative AI by taking action rather than only producing content, including setting subgoals and running multi-step workflows.
- Agent types range from simple reflex to model-based, goal-based, utility-based, and learning agents; match capability to task complexity and data maturity.

- Strengths include scale, consistency, and optimization, while risks include hallucination, opaque logic, bias, tool failures, and misalignment; human oversight and guardrails are essential.

AI in Action: Case studies show real impact, from Puma's agent-led creative pipeline and Vera Bradley's Copilot-enabled service to Lush's brand protection agents and Walmart's retail super agents.

Think Deeper:

1. What criteria should trigger human review versus full autonomy, such as customer harm risk, regulatory exposure, or brand impact?
2. How can teams balance ReAct and ReWOO within one workflow, and detect when to switch methods as uncertainty changes?
3. Which metrics best capture long-term value of agents beyond short-term lift, including trust, lifetime value, and error cost avoided?

This is where you can add appendices or other back matter.