




FALL 2026 EDITION
UPDATED TWICE WEEKLY



AI in Marketing

The Living Textbook

Dr. Ayse Ozturk

Assistant Professor of Marketing at the Darla Moore School of Business, University of South Carolina

AI IN MARKETING

THE LIVING TEXTBOOK · FALL 2026 EDITION

CHAPTER 1

Foundations of AI and Marketing

What AI is, how the tools work under the hood, where they fail,
and why they change the economics of marketing's predictions.

Dr. Ayse Ozturk

Assistant Professor of Marketing
Darla Moore School of Business
University of South Carolina

FREE SAMPLE CHAPTER



Read the living version,
updated every Monday and Wednesday
[drayseozturk.org/ai-in-marketing-textbook/
chapter-1-foundations](https://drayseozturk.org/ai-in-marketing-textbook/chapter-1-foundations)

ABOUT THIS SAMPLE

A textbook built for a field that changes every week

Most marketing textbooks treat AI as a closing chapter. *AI in Marketing: The Living Textbook* puts it at the center and keeps it current. It was written for MKTG 725, the MBA course AI in Marketing at the Darla Moore School of Business, University of South Carolina, and lives online, where a “Latest updates” box in every chapter adds new research, rulings, and launches twice a week.

Real cases

Every chapter opens with a recent case from well-known brands.

Every claim sourced

Facts, numbers, and quotes link to the original source, with a full APA reference list. Company-reported figures are labeled.

Practice built in

Key terms, a ten-question quiz, a judgment exercise, and discussion questions with model answers.

Always current

The online edition is updated every Monday and Wednesday. This PDF is a snapshot from October 4, 2026.

The seven chapters

01 Foundations of AI and Marketing **THIS SAMPLE**

What AI is, how large language models work, the jagged frontier, and why marketing has always been a prediction business.

Opening case: Would you like AI with that? Inside the drive-thru test

02 Generative AI and Content Marketing

What generative AI can make, how to prompt like a creative director, and why “made with AI” can cost a brand trust.

Opening case: Made by machines, judged by people

03 Personalization, SEO to GEO, and Retailing

Personalization at scale, the shift from search rankings to AI answers, and agentic commerce in retail.

Opening case: Amazon shut the door. Shopify opened one.

04 Marketing Automation, Workflows, and Agents

From single prompts to workflows to AI agents, the connectors behind them, and where people must stay in charge.

Opening case: The agent that did 700 jobs, and the one that gave away a PlayStation

05 Marketing Research, Segmentation, and Predictive Analytics

Predictive models, AI segmentation, synthetic data, and digital twins, and how to tell a good prediction from a flattering one.

Opening case: The focus group gets a digital twin, then an audit

06 Strategy, Implementation, Ethics, and Future Trends

Turning pilots into value, governing AI, and managing hallucinations, bias, copyright, and privacy, plus the skills that matter next.

Opening case: When AI gets it wrong, the brand pays

07 Review and Term Project: The AI Marketing Audit

One framework for the whole course and a step-by-step guide to auditing a real company’s marketing with AI.

Opening case: Strategy is what you say, too: Duolingo’s “AI-first” memo

Chapter 1 is free to read online. Chapters 2–7 are available to students enrolled in MKTG 725. Questions, feedback, or interest in using this textbook in your course: ayse.ozturk@moore.sc.edu.

Foundations of AI and Marketing

Marketing has always been a prediction business: who will buy, what to show them, when to reach out. This chapter explains what AI is, how the tools you will use all term work under the hood, and why they change the economics of those predictions.

About 40 minutes to read, plus videos and practice · Snapshot of October 4, 2026

IN THIS CHAPTER

Opening case: Would you like AI with that?

1.1 What AI is, and what it isn't

1.2 How a large language model works

1.3 The jagged frontier

1.4 Why AI reshapes marketing

1.5 Where things stand in 2026

1.6 Limits and risks to keep in view

Key takeaways

Study and practice

References

Answer key

Would you like AI with that? Inside the drive-thru test

McDonald's and Taco Bell are betting that AI can take your order. Their bumpy road there shows what AI does well, where it stumbles, and why marketers should care.

~100

McDonald's restaurants in its IBM voice-ordering test (2021–2024)

890+

Taco Bell restaurants using voice AI by July 2026

90%+

Archy's order accuracy, as claimed by McDonald's (unaudited)

On September 23, 2026, McDonald's told investors it will begin rolling out ArchIQ, a generative AI “restaurant operating system,” in 2027 ([Mira & Stone, 2026](#)). Its most visible piece is **Archy**, a voice agent built with Google that takes drive-thru orders in English and Spanish. McDonald's says Archy is more than 90% accurate and will free at least 50 labor hours per restaurant each week. Those are the company's own, unaudited figures ([PYMNTS, 2026](#)).

It is McDonald's second try. A voice-ordering test with IBM, begun in 2021, ended in June 2024 at about 100 restaurants after customers posted TikTok videos of the system taking orders from the wrong cars, multiplying orders, and producing combinations such as ice cream with ketchup and butter ([Al Jazeera, 2024](#)).

Taco Bell hit the same wall. With voice AI in roughly 500 drive-thrus, it slowed down in 2025 and rethought its rollout after complaints and pranks went viral, including one customer's order for 18,000 cups of water ([Littman, 2026](#); [Okeke, 2025](#)). Its chief digital and technology officer, Dane Mathews, told *The Wall Street*

TIMELINE

2021 · McDonald's starts testing voice ordering with IBM.

June 2024 · McDonald's ends the IBM test at about 100 restaurants after order mix-ups go viral on TikTok.

Aug 2025 · Taco Bell slows its rollout after complaints and pranks, including an order for 18,000 cups of water.

July 2026 · Taco Bell's voice AI is taking orders at more than 890 restaurants in 38 states.

Sept 2026 · McDonald's unveils ArchIQ and its voice agent, Archy, for a U.S. rollout starting in 2027.

Journal, “sometimes it lets me down, but sometimes it really surprises me” (as quoted in [Okeke, 2025](#)). Taco Bell kept going. By July 2026 its voice AI was taking orders at more than 890 restaurants, with transaction times the company describes as on par with, and in some cases faster than, traditional ordering ([Revell, 2026](#)).

WHY THIS CASE OPENS THE CHAPTER

- **Fluent is not the same as accurate** (Section 1.2). A voice agent sounds just as confident when it mishears an order.
- **The jagged frontier** (Section 1.3). Routine orders sit inside it; noise, unusual requests, and pranks can fall outside it.
- **Adoption versus value** (Section 1.5). McDonald’s makes its case in accuracy and labor hours, not in “AI” for its own sake.

Keep this question in mind: If an AI takes 9 out of 10 orders correctly, is that good enough? Who notices the tenth order, and what does it cost the brand? By the end of this chapter, you will have the tools to answer.

Sources: Al Jazeera (2024), Okeke (2025), Littman (2026), Revell (2026), Mira & Stone (2026), and PYMNTS (2026). Full citations are in the [References](#).

By the end of this chapter, you should be able to:

- Explain how artificial intelligence, machine learning, deep learning, foundation models, and generative AI nest inside one another, and why it matters which layer a vendor is selling (1.1).
- Distinguish supervised, unsupervised, and reinforcement learning, and match each to a marketing use such as churn prediction, segmentation, or send-time testing (1.1, 1.4).
- Explain how a large language model produces text by predicting one token at a time, and why its fluent output can still be false or tilted toward what the user wants to hear (1.2).
- Apply the jagged frontier to decide which marketing tasks to hand to AI and which outputs to verify closely (1.3).
- Classify a marketing AI application by the job it does (predict, recommend, generate, or act) and by Davenport, Guha, and Grewal's two dimensions: intelligence level and stand-alone versus integrated (1.4).
- Evaluate the gap between AI adoption and profit impact in marketing, and identify the five risks to keep in view when using AI: hallucinations, bias, sycophancy, privacy breaches, and workslop (1.5, 1.6).

KEY TERMS

Machine learning · Deep learning · Foundation model · Large language model (LLM) · Token · Transformer · Temperature · Hallucination · Sycophancy · Jagged frontier · Moravec's paradox · Agentic AI · Agent washing · Workslop

Definitions and examples: [AI Dictionary](#). Practice them with the flashcards at the end of this chapter.

Try this before you read

You are a marketing research associate for a new energy drink brand. Ask an AI tool (ChatGPT, Claude, or Gemini) for competitive insights on Red Bull versus Monster. Write one strong prompt and save the output. As you read this chapter, mark which parts of the answer you would trust as is and which you would verify first. (*Adapted from the Week 1 class opener.*)

1.1 What AI is, and what it isn't

Artificial intelligence is the science of getting computers to do things that normally require human intelligence: learning, reasoning, solving problems, and understanding language. It is less a single technology than a family of techniques, and the members of that family nest inside one another.

Artificial intelligence. Any technique that lets machines perform tasks we associate with human intelligence, including older rule-based expert systems.

Machine learning. Systems that learn patterns from data instead of following hand-written rules. Your email spam filter is the classic example.

Deep learning. Machine learning with many-layered neural networks. It excels at messy, unstructured data such as images, audio, and text.

Foundation models. Very large deep learning models pre-trained on broad data, then adapted to many tasks instead of being built from scratch for each one.

Large language models and generative AI. Foundation models that create new content: text, images, audio, video, and code. ChatGPT, Claude, and Gemini live here.

This nesting has two practical consequences. First, “AI” in a vendor pitch can mean anything from a 1990s rules engine to a frontier model, so always ask which layer you are buying. Overstating it has a name: *agent washing*, labeling basic automation as an autonomous AI agent (Ozturk, [AI Dictionary](#), 2026). Second, the predictive tools marketers have used for over a decade, such as churn models, propensity scores, and recommenders, are machine learning too. Generative AI did not replace them. It added a new layer on top.

Three ways machines learn

Approach	How it learns	Marketing example
Supervised	From labeled examples with known answers	Predicting which customers will cancel, trained on past customers who did and did not
Unsupervised	Finds structure in unlabeled data	Clustering customers into segments nobody defined in advance
Reinforcement	Trial and error, guided by rewards	Testing email send times or ad bids and learning which earn the most clicks

1.2 How a large language model works

At its core, a large language model (LLM) is a prediction engine. Given some text, it predicts a likely next token (a word or piece of a word), adds it, and repeats. Picture the autocomplete on your phone, trained on enormous amounts of text from websites, books, and other sources, and scaled up until it can hold a conversation (Sanderson, 2024).

Three ideas explain most of what you will see in class:

- **Transformers and attention.** The transformer architecture, introduced in 2017 by a team of researchers mostly at Google, lets a model weigh every word in a passage against every other word at once (Vaswani et al., 2017). That is how it keeps track of what “it” refers to three sentences back.
- **Pre-training, then fine-tuning.** Models first learn general language patterns from huge text collections. Further training, including feedback from human reviewers, then teaches them to follow instructions and respond helpfully (Ouyang et al., 2022).
- **Temperature.** A setting that controls how random the word choices are. Higher temperature gives more varied, more creative output; lower temperature gives more predictable output (Ozturk, AI Dictionary, 2026).



VIDEO · SCAN THE CODE OR CLICK THE TITLE

Large Language Models explained briefly (3Blue1Brown)

Grant Sanderson (3Blue1Brown), *Large Language Models Explained Briefly*, made for a Computer History Museum exhibit on chatbots (Sanderson, 2024).
youtube.com/watch?v=LPZh9BOjkQs


Why fluent is not the same as accurate

Because an LLM predicts plausible text, it can produce a confident, well-written answer that is false. This is a **hallucination**. The model is not lying; it simply has no built-in way of knowing what is true. Models also tend to tell users what they want to hear, a pattern called **sycophancy**, which researchers trace in part to training that rewards the answers people prefer (Sharma et al., 2024). A leading prompt tends to get a leading answer. Treat AI output the way you would treat a bright intern's first draft: fast, useful, and in need of checking.

1.3 The jagged frontier

AI's abilities do not rise evenly. A model can draft a sharp positioning statement in seconds, then stumble on a task that looks just as easy. Researchers call this uneven boundary the **jagged technological frontier** (Dell'Acqua et al., 2023).

Their field experiment involved 758 Boston Consulting Group consultants. On 18 realistic consulting tasks that fell inside the frontier, consultants with access to GPT-4 outperformed a control group without it. On one task deliberately chosen to fall outside the frontier, they did worse ([Dell'Acqua et al., 2023](#)):



Figures as reported in the paper's abstract. The full paper breaks results down further, by the two AI conditions tested (with and without a prompt-engineering overview).

The lesson for marketers is not that AI is weak. It is that AI fails in places you do not expect, and it sounds just as confident when it does. A related and much older idea, **Moravec's paradox**, observes that what is hard for people (chess, calculus) is often easy for machines, while what is easy for people (walking, grasping a cup) is often hard for them ([Moravec, 1988](#)). The frontier also moves: when one bottleneck breaks, such as a jump in image-generation quality, whole categories of work can suddenly come within reach ([Ozturk, AI Notes to Remember, 2026](#), summarizing Ethan Mollick). To see where the edge sits this month, check the Frontier Models page ([Ozturk, Frontier Models, 2026](#)).

Usually inside the frontier

- First drafts of copy, and many variations of it
- Summarizing reviews, transcripts, and reports
- Brainstorming segments, personas, and campaign ideas
- Turning messy notes into a structured brief

Check carefully

- Specific numbers, market sizes, and statistics
- Citations and quotes
- Recent events the model may not know about
- Legal, regulatory, and pricing claims

These two lists are classroom rules of thumb as of fall 2026, not findings from the study. The frontier shifts with every model release.

1.4 Why AI reshapes marketing

AI lowers the cost and raises the speed of the predictions marketing depends on. Generative AI adds something new on top: the cost of producing content falls sharply. Most marketing AI does one of four jobs:

Predict

Supervised models forecast churn, customer lifetime value, and likelihood to buy, so retention and acquisition budgets go where they matter most.

Recommend

Recommender systems decide what each customer sees first. Writing about Netflix in 2015, Carlos Gomez-Uribe and Neil Hunt estimated that “the combined effect of personalization and recommendations save us more than \$1B per year” ([Gomez-Uribe & Hunt, 2015](#)).

Generate

Generative models draft copy, images, audio, and video, then adapt them for each audience and channel. Chapter 2 covers this in depth.

Act

Agentic AI strings these together: it reads signals, plans a campaign, carries out steps across tools, and reports back, with people approving the key decisions. Chapter 4 covers this in depth.

A framework for classifying any marketing AI

Davenport, Guha, and Grewal sort marketing AI along two dimensions: its intelligence level (simple task automation versus machine learning) and whether it stands alone or is part of a broader platform. They describe simple stand-alone task-automation apps as a good place to start, and integrated machine-learning applications as having the greatest potential to create value ([Davenport et al., 2021](#)).

Intelligence level	Stand-alone	Integrated into a platform
Task automation (follows rules)	A website chatbot that answers a fixed set of FAQs	Automated emails triggered inside your CRM when a customer abandons a cart
Machine learning (learns from data)	A separate tool that scores leads from a spreadsheet you upload	Product recommendations built into the website and app, updated with every click

Framework from Davenport et al. (2021). The examples in the cells are ours, for illustration. See also HBR’s graphic, [The Four Kinds of Marketing AI](#) (Harvard Business Review, 2021).

1.5 Where things stand in 2026

88%

of organizations used AI in at least one business function in 2025, up from 78% in 2024. [Stanford HAI \(2026\)](#), reporting McKinsey survey data.

37%

of survey respondents attribute at least some EBIT (operating profit) impact to AI, about the same share as a year earlier. [Tinkoff et al. \(2026\)](#).

40%

of respondents at large companies (over \$1 billion in revenue) report scaling AI agents, up from 27% a year earlier. [Tinkoff et al. \(2026\)](#).

Two patterns matter most for marketers. First, McKinsey's 2026 survey finds that revenue gains are most often attributed to AI use in marketing and sales ([Tinkoff et al., 2026](#)). Stanford's AI Index, drawing on McKinsey's 2025 data, also ranks marketing and sales first for revenue gains, and summarizes studies reporting productivity gains of 14% to 15% in customer support, 26% in software development, and 50% in marketing output ([Stanford HAI, 2026](#)). Second, use is wide but profit impact is thin: most organizations use AI somewhere, yet just over a third of respondents (37%) can point to any effect on operating profit. Closing that gap between adoption and value is exactly what your AI Marketing Audit asks you to do.

1.6 Limits and risks to keep in view

- **Hallucinations.** Confident statements that are false. Verify every number, quote, and citation before it reaches a client or a slide.
- **Bias.** Models can absorb and reproduce human biases present in their training data, with consequences in targeting, pricing, and creative. IBM's overview is a good primer ([Holdsworth, 2023](#)).
- **Sycophancy.** Models tend to tell you what you want to hear ([Sharma et al., 2024](#)). Ask for counterarguments and weaknesses, not just confirmation.
- **Privacy.** Never paste customer data, personal information, or confidential plans into a consumer AI tool. This is also course policy (see the syllabus).
- **Workslop.** AI-generated work that looks polished but lacks real substance, wasting the time of whoever receives it ([Niederhoffer et al., 2025](#)).

Key takeaways

- **1.1** AI is a family of nested techniques, from rule-based systems through machine learning, deep learning, and foundation models to LLMs and generative AI, so ask which layer a “powered by AI” tool really uses. Generative AI did not replace predictive tools such as churn models and recommenders; it added a layer on top.
- **1.2** An LLM writes by predicting one token at a time, built on the 2017 transformer architecture, pre-training on huge text collections, and fine-tuning with human feedback. Because it produces plausible text rather than checked facts, it can hallucinate or tell users what they want to hear, so treat its output as a bright intern’s first draft.
- **1.3** AI’s abilities are uneven. Among 758 BCG consultants, GPT-4 users completed 12.2% more tasks, 25.1% faster, with more than 40% higher quality inside the frontier, but were 19 percentage points less likely to be correct on a task outside it. Hand off drafting and summarizing; check numbers, citations, recent events, and legal claims.
- **1.4** AI makes marketing predictions cheaper and faster, generative AI sharply cuts the cost of content, and most marketing AI does one of four jobs: predict, recommend, generate, or act. Davenport, Guha, and Grewal see stand-alone task automation as a good start and integrated machine learning as having the greatest potential to create value.
- **1.5** AI use is wide (88% of organizations used it in at least one function in 2025), yet only 37% of survey respondents attribute any EBIT impact to it. Revenue gains are most often attributed to AI in marketing and sales; closing the gap between adoption and value is the job of your AI Marketing Audit.
- **1.6** Keep five risks in view: hallucinations, bias absorbed from training data, sycophancy, privacy breaches, and workslop that looks polished but lacks substance. Verify every number, quote, and citation, ask for counterarguments, and never paste customer data or confidential plans into a consumer AI tool.

Watch

Two full-length public lectures from The Turing Lectures series, filmed at the Royal Institution in partnership with The Alan Turing Institute.

VIDEO · SCAN THE CODE OR CLICK THE TITLE



What is generative AI and how does it work? The Turing Lectures with Mirella Lapata

Mirella Lapata, “What is generative AI and how does it work?” (The Royal Institution, 2023a).

youtube.com/watch?v=_6R7Ym6Vy_I

VIDEO · SCAN THE CODE OR CLICK THE TITLE



What's the future for generative AI? The Turing Lectures with Mike Wooldridge

Mike Wooldridge, “What's the future for generative AI?” (The Royal Institution, 2023b).

youtube.com/watch?v=b76gsOSkHB4

Study and practice

Four ways to lock in this chapter, from quick recall to applied judgment. Answers are in the Answer key at the end of the chapter. The online chapter has interactive versions of the flashcards, quiz, and sorting exercise.

Key terms to know

Cover the definition with your hand, say it out loud, then check.

Machine learning

Systems that learn patterns from data instead of following hand-written rules. *Example: an email spam filter.*

Deep learning

Machine learning that uses many-layered neural networks; strong with images, audio, and text.

Foundation model

A very large model pre-trained on broad data, then adapted to many different tasks.

Large language model (LLM)

A foundation model that generates text by repeatedly predicting the next token. *Examples: the models behind ChatGPT, Claude, and Gemini.*

Token

A word or piece of a word: the unit an LLM reads and predicts.

Transformer

The 2017 model design that uses attention to weigh every word in a passage against every other word. *Vaswani et al., 2017*

Temperature

A setting that controls how random the output is. Higher means more varied and creative; lower means more predictable.

Hallucination

A fluent, confident output that is false. *Verify numbers, quotes, and citations.*

Sycophancy

The tendency of AI assistants to tell users what they want to hear. *Sharma et al., 2024*

Supervised vs. unsupervised learning

Supervised learns from labeled examples (e.g., churn prediction). Unsupervised finds structure in unlabeled data (e.g., clustering segments).

Jagged frontier

The uneven edge of AI ability: excellent at some tasks, unexpectedly poor at others that look similar. *Dell'Acqua et al., 2023*

Moravec's paradox

What is hard for people (chess, calculus) is often easy for machines, and what is easy for people (walking, grasping) is often hard. *Moravec, 1988*

Agentic AI

AI that plans and carries out multi-step actions toward a goal with limited human supervision.

Agent washing

Marketing basic automation or rule-based tools as if they were autonomous AI agents. *Ozturk, AI Dictionary, 2026*

Workslop

AI-generated work that looks polished but lacks real substance. *Niederhoffer et al., 2025*

Practice quiz

Ten multiple-choice questions, like the in-class quiz. Choose one answer for each.

- 1. Which statement best describes how a large language model generates text?**
 - A. It looks up answers in a database of verified facts.
 - B. It copies the closest matching web page.
 - C. It runs a separate hand-written rule for each question.
 - D. It predicts a likely next token, one token at a time.
- 2. A model trained on past customers labeled “stayed” or “canceled” to predict who will cancel next is an example of:**
 - A. Unsupervised learning
 - B. Supervised learning
 - C. Reinforcement learning
 - D. Agent washing
- 3. Grouping customers into segments that nobody defined in advance uses:**
 - A. Unsupervised learning
 - B. Supervised learning
 - C. Reinforcement learning
 - D. A transformer
- 4. In Dell'Acqua et al. (2023), what happened on the task deliberately placed outside AI's frontier?**
 - A. Consultants using AI finished faster with the same accuracy.
 - B. There was no difference between groups.
 - C. Consultants using AI were 19 percentage points less likely to produce a correct solution.
 - D. Consultants refused to use AI on that task.

5. **According to Davenport et al. (2021), which kind of marketing AI has the greatest potential to create value?**
- A. Stand-alone task automation
 - B. Integrated task automation
 - C. Integrated machine learning
 - D. Stand-alone machine learning
6. **Raising an LLM's temperature setting generally makes its output:**
- A. More predictable
 - B. More varied and creative
 - C. More accurate
 - D. Shorter
7. **In McKinsey's 2026 State of AI survey (Tinkoff et al., 2026), what share of respondents attribute at least some EBIT impact to AI?**
- A. 12%
 - B. 63%
 - C. 88%
 - D. 37%
8. **Why can an LLM write a polished paragraph containing a made-up statistic?**
- A. It is designed to deceive users.
 - B. Statistics are stored at random in its memory.
 - C. It generates plausible text and has no built-in way of checking what is true.
 - D. It only makes mistakes with numbers.
9. **A vendor calls its tool an "AI agent," but it only sends the same fixed reply whenever an email arrives. This is an example of:**
- A. Agent washing
 - B. Agentic AI
 - C. Reinforcement learning
 - D. Hallucination
10. **In the nested picture of AI, which layer is the narrowest?**
- A. Artificial intelligence
 - B. Large language models and generative AI
 - C. Machine learning
 - D. Deep learning

Inside or outside the frontier?

For each marketing task, decide: is AI usually reliable here, or should you check its work carefully? These are rules of thumb as of fall 2026, not hard rules.

Marketing task	Your call	
Stating the current market size of the U.S. energy drink category in a client deck	☐ Usually reliable	☐ Check carefully
Turning messy meeting notes into a structured creative brief	☐ Usually reliable	☐ Check carefully
Summarizing 500 customer reviews into recurring themes	☐ Usually reliable	☐ Check carefully
Confirming that a promotional claim complies with advertising law	☐ Usually reliable	☐ Check carefully
Reporting what a competitor announced last week	☐ Usually reliable	☐ Check carefully
Drafting ten subject-line variations for an email campaign	☐ Usually reliable	☐ Check carefully
Brainstorming personas for a new product launch	☐ Usually reliable	☐ Check carefully
Citing an academic study to support a pricing recommendation	☐ Usually reliable	☐ Check carefully

Check your understanding

Short-answer and discussion questions. Write your answer first, then compare it with the model answer in the Answer key.

1. A vendor says its email tool is “powered by AI.” What two questions would you ask to find out what you are actually buying?

2. Explain to a skeptical manager why an LLM can produce a polished paragraph that contains a made-up statistic.

3. In the BCG experiment, consultants using AI did better on most tasks but worse on one. What was different about that task, and what does it mean for your work?

4. Discussion: where would ChatGPT fit in the Davenport, Guha, and Grewal framework? Does generative AI fit the framework, or strain it?

5. Return to the opening case. McDonald's says Archy is more than 90% accurate. Using the jagged frontier, which drive-thru orders would you expect the AI to handle well, which might trip it up, and what safeguards would you recommend?

APPLY IT · TERM PROJECT

Start your AI Marketing Audit

Pick a company you might audit for your term project. List five of its marketing activities across the customer funnel: awareness, consideration, purchase, retention, and advocacy. For each one, note (1) whether it likely uses AI today, (2) which kind (rules-based task automation, or AI that predicts, recommends, generates, or acts), and (3) whether the task sits inside or outside the jagged frontier. Keep this table; you will build on it in later chapters.

FEEDBACK

Help improve this chapter

Was this chapter useful? Did you spot an error, a broken link, or something unclear? Tell Dr. Ozturk in a minute. Each button opens a short, pre-filled email.

Yes, it helped

Report an error or broken link

Suggest an improvement

Latest updates as of October 4, 2026

The online chapter adds new developments every Monday and Wednesday. These were the newest items when this edition was printed; scan the code on the cover for the current list.

- **Oct 2, 2026.** McKinsey's latest State of AI survey finds that 44% of respondents now say AI is scaling across their enterprise, up from 38% a year earlier, while the share attributing at least some EBIT impact to AI (37%) is essentially unchanged. Revenue gains are still most often attributed to AI in marketing and sales. Connects to 1.5. [McKinsey, Aug 25, 2026](#) (mckinsey.com)
- **Oct 2, 2026.** Stanford's 2026 AI Index reports that global corporate investment in AI, counting mergers and acquisitions, minority stakes, private investment, and public offerings, reached \$581.69 billion in 2025, a 129.9% increase over the prior year. Connects to 1.5. [Stanford HAI, 2026](#) (hai.stanford.edu)

References

APA style (7th edition). Dr. Ozturk's own resource pages are cited by page name, for example (Ozturk, AI Dictionary, 2026), because they are updated continuously. Some publications require a subscription; USC Libraries can often provide access. Sources for "Latest updates" items are linked within each item.

- Al Jazeera. (2024, June 19). McDonald's scraps AI pilot after order mix-ups go viral. *Al Jazeera*.
<https://www.aljazeera.com/economy/2024/6/19/mcdonalds-scrap-ai-pilot-at-drive-through-outlets-after-order-mix-ups>
- Davenport, T. H., Guha, A., & Grewal, D. (2021, July–August). How to design an AI marketing strategy. *Harvard Business Review*. <https://hbr.org/2021/07/how-to-design-an-ai-marketing-strategy>
- Dell'Acqua, F., McFowland, E., III, Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013; forthcoming in *Organization Science*). SSRN. <https://doi.org/10.2139/ssrn.4573321>
- Gomez-Urbe, C. A., & Hunt, N. (2015). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4), Article 13. <https://doi.org/10.1145/2843948>
- Harvard Business Review. (2021, July). *The four kinds of marketing AI* [Data visualization]. <https://hbr.org/data-visuals/2021/07/the-four-kinds-of-marketing-ai>
- Holdsworth, J. (2023, December 22). What is AI bias? IBM. <https://www.ibm.com/think/topics/ai-bias>
- Littman, J. (2026, July 7). Taco Bell revs up drive-thru AI deployment. *Restaurant Dive*.
<https://www.restaurantdive.com/news/taco-bell-omilia-drive-thru-ai-deployment/824564/>
- Mira, L., & Stone, D. (2026, September 23). McDonald's commits \$8.5 billion to restaurant modernization as ArchIQ moves AI into daily operations. *Restaurant Technology News*.
<https://restauranttechnologynews.com/2026/09/mcdonalds-commits-8-5-billion-to-restaurant-modernization-as-archiq-moves-ai-into-daily-operations/>
- Moravec, H. (1988). *Mind children: The future of robot and human intelligence*. Harvard University Press.
<https://www.hup.harvard.edu/books/9780674576186>
- Niederhoffer, K., Rosen Kellerman, G., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (2025, September 22). AI-generated "workslop" is destroying productivity. *Harvard Business Review*. <https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity>
- Okeke, F. (2025, September 1). Taco Bell scales back AI tests after customer complaints. *Computing*.
<https://www.computing.co.uk/news/2025/ai/taco-bell-scales-back-ai-tests-after-customer-complaints>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., et al. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35. <https://arxiv.org/abs/2203.02155>
- Ozturk, A. (2026). *AI dictionary* [Web page, updated continuously; last updated September 28, 2026]. [drayseozturk.org. https://drayseozturk.org/ai-dictionary/](https://drayseozturk.org/ai-dictionary/)
- Ozturk, A. (2026). *AI notes to remember* [Web page, updated continuously; last updated September 30, 2026]. [drayseozturk.org. https://drayseozturk.org/ai-notes/](https://drayseozturk.org/ai-notes/)
- Ozturk, A. (2026). *Frontier models* [Web page, updated continuously; last updated October 2, 2026]. [drayseozturk.org. https://drayseozturk.org/frontier_models/](https://drayseozturk.org/frontier_models/)
- PYMNTS. (2026, September 29). McDonald's puts AI in charge of the drive-thru. PYMNTS.
<https://www.pymnts.com/news/artificial-intelligence/2026/mcdonalds-puts-ai-in-charge-of-the-drive-thru/>
- Revell, E. (2026, July 8). Taco Bell ramps up voice AI use across nearly 900 drive-thrus. *Fox Business*.
<https://www.foxbusiness.com/lifestyle/taco-bell-ramps-up-voice-ai-use-across-drive-thrus>

- The Royal Institution. (2023a, October 12). *What is generative AI and how does it work? – The Turing Lectures with Mirella Lapata* [Video]. YouTube. https://www.youtube.com/watch?v=_6R7Ym6Vy_I
- The Royal Institution. (2023b, December 19). *What's the future for generative AI? – The Turing Lectures with Mike Wooldridge* [Video]. YouTube. <https://www.youtube.com/watch?v=b76gsOSkHB4>
- Sanderson, G. (2024, November 20). *Large language models explained briefly* [Video lesson]. 3Blue1Brown. <https://www.3blue1brown.com/lessons/mini-llm>
- Sharma, M., Tong, M., Korbak, T., et al. (2024). Towards understanding sycophancy in language models. *International Conference on Learning Representations (ICLR 2024)*. <https://arxiv.org/abs/2310.13548>
- Stanford Institute for Human-Centered Artificial Intelligence [Stanford HAI]. (2026). *The 2026 AI Index report*. Stanford University. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
- Tinkoff, D., Van der Veken, L., Chui, M., & Balakrishnan, T. (2026, August 25). *The state of AI in 2026: On the road to ROI*. McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30. <https://arxiv.org/abs/1706.03762>

Go deeper (optional)

- Westerman, G., & Webster, M. (2025, January 22). Generate value from GenAI with “small t” transformations. *MIT Sloan Management Review*. <https://sloanreview.mit.edu/article/generate-value-from-gen-ai-with-small-t-transformations/> Why modest, focused AI projects often pay off first.
- Course books by Dr. Ozturk (on Blackboard): *Applied AI in Marketing* and *Artificial Intelligence Foundations and Literacy*, plus her password-protected course page for enrolled students, [Course: “AI in Marketing”](#).

HOW TO CITE THIS CHAPTER

If you use ideas, figures, or exercises from this chapter in your teaching, research, or writing, please cite it as shown below, with the date you accessed it, because the chapter is updated regularly. In-text citation: (Ozturk, 2026). For a specific fact, statistic, or quotation, cite the original source listed in the References.

APA 7

Ozturk, A. (2026). Foundations of AI and marketing (Chapter 1). In *AI in marketing: The living textbook*. Retrieved [Month Day, Year], from <https://drayseozturk.org/ai-in-marketing-textbook/chapter-1-foundations/>

Answer key

Practice quiz

1. **D.** It predicts a likely next token, one token at a time.
An LLM is a prediction engine: it predicts a plausible next token, adds it, and repeats (Section 1.2).
2. **B.** Supervised learning
Learning from labeled examples with known answers is supervised learning (Section 1.1).
3. **A.** Unsupervised learning
Finding structure in unlabeled data, such as clusters, is unsupervised learning (Section 1.1).
4. **C.** Consultants using AI were 19 percentage points less likely to produce a correct solution.
Outside the frontier, consultants using AI were 19 percentage points less likely to get it right than those without AI (Section 1.3).
5. **C.** Integrated machine learning
They call stand-alone task automation a good place to start, but see integrated machine learning as having the greatest value potential (Section 1.4).
6. **B.** More varied and creative
Higher temperature adds randomness to word choice: more variety, less predictability (Section 1.2).
7. **D.** 37%
37%, about the same as the prior year. The 88% figure is the share of organizations using AI in at least one function in 2025 (Section 1.5).
8. **C.** It generates plausible text and has no built-in way of checking what is true.
Fluency is not accuracy: plausible-sounding numbers are exactly what a next-token predictor can produce (Section 1.2).
9. **A.** Agent washing
Labeling simple rule-based automation as an autonomous agent is agent washing (Section 1.1).
10. **B.** Large language models and generative AI
LLMs and generative AI sit inside foundation models, inside deep learning, inside machine learning, inside AI (Section 1.1).

Inside or outside the frontier?

1. **Check carefully.** Stating the current market size of the U.S. energy drink category in a client deck.
Specific numbers are where models can invent plausible figures. Find the original source.
2. **Usually reliable.** Turning messy meeting notes into a structured creative brief.
Restructuring your own material is usually reliable, because the facts come from you.
3. **Usually reliable.** Summarizing 500 customer reviews into recurring themes.
Summarizing text is usually reliable. Spot-check a few reviews against the themes.
4. **Check carefully.** Confirming that a promotional claim complies with advertising law.
Legal and regulatory questions need a qualified human and the actual rules.
5. **Check carefully.** Reporting what a competitor announced last week.
Recent events may be missing from the model's training data. Check a dated source.

6. **Usually reliable.** Drafting ten subject-line variations for an email campaign.
Generating variations is a classic strength. You still pick and edit the winners.
7. **Usually reliable.** Brainstorming personas for a new product launch.
Brainstorming is low risk and high value. Validate the personas with real data later.
8. **Check carefully.** Citing an academic study to support a pricing recommendation.
Models can produce realistic-looking citations that do not exist. Open and read the source.

Check your understanding: model answers

1. **A vendor says its email tool is “powered by AI.” What two questions would you ask to find out what you are actually buying?**
Does it learn from data (machine learning), or does it follow fixed rules (task automation)? And does it stand alone, or does it plug into your CRM and customer data? You might also ask which model it uses and what data it learns from.
2. **Explain to a skeptical manager why an LLM can produce a polished paragraph that contains a made-up statistic.**
It generates the most plausible next words, not verified facts, and a realistic-sounding number is plausible text. Fluency tells you nothing about accuracy, so numbers and citations always need checking against a source.
3. **In the BCG experiment, consultants using AI did better on most tasks but worse on one. What was different about that task, and what does it mean for your work?**
That task was designed to sit outside the frontier of what the AI could do well, yet its output still looked convincing. The takeaway: learn where the frontier lies for your own tasks, and check AI output most carefully exactly where it feels most convincing.
4. **Discussion: where would ChatGPT fit in the Davenport, Guha, and Grewal framework? Does generative AI fit the framework, or strain it?**
There is no single right answer. Used on its own in a browser, ChatGPT works like stand-alone machine learning; the same kind of model built into a CRM or ad platform is integrated machine learning. Some argue the framework still works; others argue generative AI adds a dimension (creating content, not just predicting) that a 2021 framework could not fully anticipate. Be ready to defend your view in class.
5. **Return to the opening case. McDonald’s says Archy is more than 90% accurate. Using the jagged frontier, which drive-thru orders would you expect the AI to handle well, which might trip it up, and what safeguards would you recommend?**
Routine, clearly spoken orders from the standard menu sit inside the frontier. Background noise, heavy customization, unusual requests, and deliberate pranks (like 18,000 cups of water) can fall outside it. Sensible safeguards include limits on order size, an on-screen order confirmation, and a crew member ready to take over: a human in the loop. Also note that the 90% figure is the company’s own and has not been independently audited.

ABOUT THE AUTHOR

Dr. Ayse Ozturk

Dr. Ayse Ozturk is an Assistant Professor of Marketing at the Darla Moore School of Business, University of South Carolina. Before academia she worked at PricewaterhouseCoopers, Deloitte, and Peugeot. Her research on artificial intelligence, marketing strategy, international marketing, and sustainability appears in the *Journal of International Business Studies*, *International Business Review*, and *Thunderbird International Business Review*. She designed and teaches the MBA course AI in Marketing and leads AI workshops for faculty and professionals in the U.S. and Canada.

Ph.D. in Marketing, Georgia State University

Provost's AI Teaching Fellow

Michael J. Mungo Undergraduate Teaching Award, 2026

Featured in *The Wall Street Journal*

Featured in OpenAI Academy's Global Faculty AI Project

HOW THIS TEXTBOOK IS BUILT

SOURCING

Peer-reviewed research, university studies, leading business publications, and regulators come first; trade press only when no stronger source exists. Company figures are labeled as company-reported.

AI USE

Chapters are drafted with the help of AI tools and reviewed, edited, and approved by Dr. Ozturk. Articles are summarized and linked, never reproduced.

VERIFICATION

Every number, quote, author, and date is checked against the original source before it appears. If something cannot be verified, it is left out.

CORRECTIONS

Spot an error or a broken link? Email ayse.ozturk@moore.sc.edu.



Keep reading

The full textbook, the AI Dictionary, and free AI-literacy resources live at drayseozturk.org.

Contact: ayse.ozturk@moore.sc.edu