

The Business Case for Analytics and Artificial Intelligence

In this chapter, we will examine the following:

- Modern AI—evolution and competitive advantage
- The democratization of AI and competitive advantage
- Real-world applications of AI and generative AI
- Analytics and AI investment trends
- Justifying investments with concrete outcomes

Five years from now, the leaders who succeeded in this moment may say they made the right call on artificial intelligence (AI). The others may talk about the “hype cycle” and wonder where their advantage went. We are at a generational inflection point. AI is being embedded in daily workflows and reshaping how we compete. From predictive analytics to generative and agentic AI, organizations are discovering new ways to cut costs, boost productivity, and open new markets.

But chasing hype is not a strategy. Too many leadership teams are trying to pilot chatbots, co-pilots, and generative tools without the solid data foundations, governance, or operational models to make them stick. The result can be proof-of-concepts that don’t scale and security and compliance risks that they didn’t anticipate. If you want AI to create measurable value, you need to start with a clear business case, one that ties investment to strategy, outcomes, and the capabilities required to deliver them.

This chapter examines the business case for AI, exploring its history, how it is evolving, and how it can create value.

Modern AI—Evolution and Competitive Advantage

Everyone is talking about AI. It is on the cover of magazines, in the news, and the subject of conferences and events. Legislative bodies, government entities, and others are trying to regulate it. It is a major focus for tech companies, which are embedding it in their software and creating new innovations using AI. I've spoken with many organizations that say they are under executive and board-level pressure to “do something” with AI.

Yet, AI is not new. The idea that machines could act intelligently has been around since the time of the ancient Greeks. In the 1950s, when John McCarthy coined the term at a summer research workshop at Dartmouth College, he described it as “making a machine behave in ways that would be called intelligent if a human were so behaving.”¹

Generally, people recognize several major waves of AI. The first wave focused on logic, rules, and symbolic reasoning. For instance, the ELIZA system, created by Joseph Weizenbaum at MIT in the mid-1960s, gave the impression of a conversation.² Known now as the first-ever chatbot, it was a rule-based natural language processing system that could simulate conversation by identifying keywords in user input and responding with predefined scripts.

The second wave began in the 1990s. I think of this period as the rise of the idea of data mining—the process of discovering patterns, trends, and useful insights from large datasets using statistical, machine learning, and analytical techniques. In other words, it marked a shift beyond traditional statistics toward the use of computer algorithms such as decision trees and basic neural networks.

This was when I first started to use AI. I remember experimenting with the C4.5 machine learning algorithm³ at Bell Labs during that time. This decision tree algorithm (what is now considered a simple machine learning algorithm) was popular because it could classify data into distinct categories with a high degree of accuracy. It also influenced the development of modern ensemble methods such as random forests and gradient boosted trees (we'll talk more about these techniques in Chapter 15), which build on the same core principles. I was impressed by the ability of C4.5 to generate rules that could, for example,

¹ McCarthy, J., Minsky, M.L., Rochester, N., and Shannon, C.E. (2006). A proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine*, 27(4), p. 12. Available at: doi.org/10.1609/aimag.v27i4.1904 [Accessed 7 June 2025].

² Weizenbaum, J. (1966). ELIZA—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), pp. 36–45.

³ Quinlan, J.R. (1993). *C4.5: Programs for machine learning*. California: Morgan Kaufman Publishers.

estimate the probability that a customer would churn. The importance of these kinds of algorithms was that they could learn patterns without rules.

Use cases for data mining included fraud detection, customer segmentation, market basket analysis, churn prediction, credit scoring, and demand forecasting. Machine learning algorithms then became more complex, supported by the compute power available to run them. Deep neural networks emerged (neural networks with multiple layers), giving rise to the “modern” AI of today.

Much of the hype around the current wave is a result of the introduction of generative AI—AI systems that can automatically create new content, such as text, images, code, or audio, by learning patterns from existing data and generating outputs. (See the following sidebar for the difference between traditional AI and generative AI.) These systems have emerged as a result of recent advances in algorithms such as the popular GPT-5. Many view these systems as revolutionary—on par with the introduction of the internet, e-commerce, and smartphones. And it is quite remarkable what they can do.

That’s a brief history of the evolution of AI to put it into context. But what about modern AI and how it can make a company more competitive? To effectively compete, a company needs a differentiated or distinct product offering that generates high customer satisfaction and retention. It needs to achieve operational efficiencies. It must have the ability to innovate and the ability to adapt to a changing landscape. Of course, it needs the financial and talent resources to accomplish this. I would argue that the ability to make evidence-based decisions and using data are also key to competitive advantage.

Using AI algorithms and building AI products can provide competitive advantage. For instance, if I can predict the probability that a customer is going to churn, I can reach out to them to try to prevent this. If I can understand the sentiment of my customers and what is causing negative sentiment, I can start to improve my products. If I can segment customers into distinct groups, I can better understand and market to them. If I can predict when my machinery needs maintenance, I can avoid excessive downtime and save money.

These examples use AI algorithms. If I can embed these AI algorithms into applications, such as recommendation engines (such as those used on Amazon), I have a better chance of increasing my company’s revenue. These recommendation engines analyze a wide range of data, including customers’ past purchases, items they’ve viewed, items in their cart, ratings, and even the behavior of similar customers. More personalized recommendations lead to potentially more purchases and greater revenue.

Numerous industry studies confirm that organizations leveraging AI are more likely to achieve measurable business impact than those that do not. For example, Accenture’s *The Art of AI Maturity* found that top performers in AI maturity (what they called AI Achievers) achieved 50% higher revenue growth

than peers.⁴ McKinsey's *State of AI* survey reported that more than a quarter of respondents attributed at least 5% of earnings before interest and taxes (EBIT) directly to AI adoption.⁵ Deloitte found that those further along in AI were significantly more likely to report revenue-generating results, such as "entering new markets/expanding services to new constituents, creating new products/programs or services, or enabling new business/service models."⁶

TRADITIONAL AI VS. GENERATIVE AI—A SIMPLE EXPLANATION

AI is an umbrella term that includes a wide range of methodologies and technologies. It draws on advances in disciplines such as mathematics, computer science, computational linguistics, cognitive sciences, and robotics. Popular AI technologies include machine learning (ML), natural language processing (NLP), and neural networks that form the foundation for the intelligent capabilities seen in modern AI systems.

Traditionally, organizations have applied AI to solve a variety of business problems, such as predicting customer churn, detecting fraud, recommending products, anticipating equipment maintenance, and forecasting disease progression. Many of these use cases rely on supervised learning, one of the most accessible and widely used approaches. In supervised learning, a model is trained on a labeled dataset—data where the outcome (or target of interest) is already known—to recognize patterns and predict future results.

For instance, a telecom company might train a model to predict customer churn using features (e.g., attributes) such as customer tenure, average monthly spend, and number of service plans. The model uses these features to learn how different customer behaviors relate to the likelihood of churn. While these patterns may be difficult for a human to detect in large datasets, machine learning algorithms such as decision trees can surface probabilistic rules—for example, identifying that customers with more than three plans and more than three years of tenure are at higher risk of leaving.

Developing traditional AI models involves a sequence of steps: collecting and preparing the data, training and tuning the model (using known outcomes—e.g., the historical data), validating and testing it for accuracy, and deploying it to production.

(continues)

⁴The art of AI maturity: Advancing from practice to performance 2021. (n.d.). Accenture. Available at: <https://www.accenture.com/content/dam/system-files/acom/custom-code/ai-maturity/Accenture-Art-of-AI-Maturity-Report-Global-Revised.pdf#zoom=40> [Accessed 7 August 2025].

⁵The State of AI in 2021. (2021). McKinsey. Available at: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/global-survey-the-state-of-ai-in-2021> [Accessed 7 August 2025].

⁶Deloitte. (2022). *Deloitte's state of AI in the enterprise*, 5th edition report. Available at: <https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2024/us-ai-institute-state-of-ai-fifth-edition.pdf> [Accessed 7 August 2025].

**TRADITIONAL AI VS. GENERATIVE AI—A SIMPLE EXPLANATION
(CONTINUED)**

Once in production, the model must be monitored over time to ensure that it continues performing well, as all models tend to degrade if left unchecked.

In generative AI, rather than predicting outcomes, systems are designed to create new outputs—such as text, images, music, or code—based on what they’ve learned from large training datasets. These systems are powered by transformer models, including large language models (LLMs) trained on a corpus that is (in many cases) the internet. The result is the ability to generate content that mimics or extends the patterns seen in their training data. Probably most of you reading this book have already used some sort of LLM such as a GPT to create an email or do some research. Instead of retrieving facts or calculating outcomes directly, these systems operate by predicting the next word in a sequence based on prior examples in their training data. As discussed later in the book, this can lead to issues and requires governance of both the data and the model.

The Democratization of AI and Competitive Advantage

I believe we’ve reached an inflection point—and an important one. With the advent of generative AI, AI is becoming consumerized and democratized in the workplace; that is, it is becoming more accessible to everyone. Think about all the ways you use AI in your daily life (e.g., virtual assistants such as Apple’s Siri or health tracking apps like Fitbit or Apple Watch). Now, many people in business are using AI as part of their daily work. It’s in their spreadsheets, emails, and meeting notes. It’s part of the everyday workflow, whether you’re in marketing, operations, or finance. Developers are using it to generate code; marketers are using it to generate campaigns. AI is no longer just the realm of data scientists and other quants. This is a major change.

This democratization is powerful—but also tricky because it leads to a key question: *How do you go beyond the surface and really use AI to compete?*

Part of the answer lies with your own data. Generative AI can do amazing things out of the box, but the real power is when you use these models with your proprietary data. Your data can help turn generic AI into a strategic asset. It can help you build smarter tools, uncover unique insights, and differentiate in ways that others can’t replicate.

Think about it. There is no doubt that organizations using analytics and AI to gain insights from their data have a competitive advantage. In my research at TDWI, I routinely look at the characteristics of those organizations that state that they have been successful in their analytics programs against those who haven’t. They are more likely to be implementing advanced analytics and particularly

AI. Additionally, they make sure to implement certain programs, such as data literacy and data governance that help them to build skills, trust insights, and compete more effectively.⁷

For example, if an organization can determine which customers are likely to churn, it can reach out to those customers with an offer that may prevent them from leaving. Ideally, the organization already has insights about the customer, including segmentation and behavioral understanding. Then, it can target these high-risk customers and use generative AI with AI-generated personalized campaigns.

Similarly, if a manufacturing company uses a predictive maintenance model (built using traditional machine learning) and identifies a system at risk of failing, it can dispatch a technician to fix the issue. That technician, while on site, might use an assistant system (developed using generative AI) that draws from manuals, diagrams, historical maintenance logs, and error code documentation to answer questions as needed and replace the part at risk.

note TDWI (www.tdwi.org) is a leading source of education and research for data and analytics professionals working across AI, business intelligence (BI), and data management. For more than 30 years, TDWI has empowered organizations with the knowledge and skills to build future-ready data and AI strategies and solutions through research-backed learning, its expert faculty, and a vibrant global community. The company provides in-person conferences, virtual events, on-demand courses, research reports, and memberships, designed to meet the evolving needs of today's data- and AI-enabled enterprises.

Real-world Applications of AI and Generative AI

There are numerous examples of how organizations are using AI together with their company data to compete. Netflix's recommendation engine uses AI to analyze customers' viewing histories and compare them with patterns from users with similar tastes to suggest movies and shows they're likely to enjoy. This creates a personalized experience. Amazon uses AI to optimize its supply chain and delivery routes. It also uses it to forecast demand based on historical information and employs techniques such as deep learning to determine whether merchandise is damaged.

The UPS ORION (On-road Integrated Optimization and Navigation) and its newer navigation tool called UPSNav system use AI to optimize delivery routes

⁷ TDWI. (2023). Achieving Success with Modern Analytics. *TDWI Best Practices Report*, p. 12. Available at: tdwi.org/research/2023/06/adv-all-best-practices-report-achieving-success-with-modern-analytics.aspx [Accessed 7 June 2025].

and provide directions in precise detail.⁸ Sports teams use AI to market to their fans, and retail companies use it to push offers to loyalty members.

There are also numerous examples of how companies are using generative AI with their company data to compete. For instance, chatbots are a very popular use case. Daily Harvest, a meal delivery service, uses AI in chatbots in its customer service channels and lets customers easily make changes to their plans, such as skipping their next order.⁹ Virtual assistant chatbots are also popular. Portfolio managers and analysts at JPMorgan use a system called Smart Monitor, which the firm estimates has helped reduce time researching a topic. It pulls in data from earnings calls, filings, and other sources, generating tailored alerts and analysis for their asset managers and wealth management professionals.¹⁰

I see many companies using generative AI on their call center notes or incident reports to summarize them and classify problems into categories. This provides valuable insights as they work toward improving customer service. Several companies I've spoken to are developing tools for their sales teams to build dossiers on potential target clients before going into meetings with them. This data may come from their customer relationship management (CRM) systems as well as external sources. All the salesperson needs to do is specify the client, and they can obtain useful intelligence about them. Some systems even generate scripts for salespeople to use, complete with the tone they should adopt for specific clients.

Agentic AI moves the needle even further forward. Although still in the early days, the idea here is to build agents that can act autonomously and proactively toward a specific goal. For instance, a supply chain agentic system might consist of a number of agents that are responsible for different tasks in the supply chain, such as inventory management or route optimization. Another example is in automating the issuance of company devices. These applications are deployed as workflows. Agents can independently retrieve information, take action (such as updating records), and coordinate across systems.

⁸ UPS. (2018). UPS deploys purpose-built navigation for UPS service personnel [Press release]. Available at: about.ups.com/us/en/newsroom/press-releases/innovation-driven/ups-deploys-purpose-built-navigation-for-ups-service-personnel.html [Accessed 8 June 2025].

⁹ Hightower, S. (2005). Daily Harvest's AI strategy is giving customers what they want. *Business Insider*. Available at: www.businessinsider.com/daily-harvest-implements-ai-to-improve-meal-delivery-customer-care-2025-2 [Accessed 8 June 2025].

¹⁰ Abrego, M. (2005). AI is core to JPMorgan's \$18 billion tech investment. Here's what its execs revealed about how it's reshaping the bank. *Business Insider*. Available at: www.businessinsider.com/jpmorgan-how-artificial-intelligence-transforming-workflows-efficiencies-2025-5 [Accessed 8 June 2025].

Table 1.1 outlines use cases currently in production for these various kinds of AI, starting with traditional machine learning (used to find patterns and make predictions), through to generative AI using company data (used to generate output), and then to agentic AI (taking action). This list is not meant to be exhaustive, but it provides some idea as to the kinds of applications organizations are building and utilizing today. On this list I've included code generation, although this is not directly related to data. This is becoming a popular use case for generative AI. A new term has even emerged, called “vibe” coding, which was coined by OpenAI cofounder Andrej Karpathy, who noted on X that it is “where you fully give into the vibes, embrace exponentials, and forget that the code even exists.”¹¹ He noted it was possible because the LLMs are getting so good.

Table 1.1: Examples of AI

TYPE OF AI	EXAMPLE USE CASES (THAT USE COMPANY DATA)
Traditional AI/ML (finds patterns and makes predictions)	Horizontal use cases
	■ Churn prediction ■ Fraud detection ■ Customer segmentation ■ Employee attrition prediction ■ Risk assessment ■ Recommendation engines ■ Demand forecasting ■ Market basket analysis ■ Dynamic pricing ■ Personalization
	Vertical use cases
	■ Patient outcome predictions ■ Predictive maintenance ■ Credit card scoring ■ Supply chain optimization ■ Route optimization ■ Network monitoring ■ Disease detection ■ Cybersecurity threat detection ■ Content moderation ■ Precision agriculture ■ Quality control and process optimization ■ Predictive operations management ■ Smart cities ■ Autonomous driving

(continues)

¹¹ Karpathy, A. (2025). [X post] 8 June. Available at: x.com/karpathy/status/1886192184808149383?lang=en [Accessed 8 June 2025].

Table 1.1: (Continued)

TYPE OF AI	EXAMPLE USE CASES (THAT USE COMPANY DATA)
Generative AI (creates new content)	Horizontal use cases ■ Customer support chatbots (external facing) ■ Employee onboarding chatbots (internal facing) ■ Sales support assistants ■ Personalized marketing campaigns ■ Summarizing call center notes, trouble tickets, incident reports ■ Maintenance assistants ■ Analyzing company data ■ Writing product descriptions at scale ■ Documenting internal compliance and workflows ■ Regulatory submission ■ Code generation ■ Customer communication drafts about incidents Vertical use cases ■ Drug discovery support ■ Summarizing physician visits ■ Creating design options for a manufactured product ■ Creating optimized supply chain replenishment schedules ■ Generating tailored investment recommendations, portfolio summaries, or financial planning reports ■ Creating detailed retail product blurbs
Agentic AI (takes action in a business context)	■ Resolving customer complaints and escalating issues ■ Managing inventory restocking ■ Monitoring facilities and alerting maintenance ■ Intelligent IT support agent ■ Maintenance agents for refineries

As previously mentioned, software vendors are embedding AI in applications. This, of course, is the nature of consumerized AI products and some of the AI applications listed in Table 1.1 and described previously. However, it is also important to note that software vendors have been infusing AI into the data and analytics life cycle to make those applications easier to use for data and analytics teams—and they’ve been doing this for years.

For instance, during data collection and profiling, AI tools are being used to automate the discovery of data quality issues, classify sensitive information such as personal identifiers, and extract metadata using intelligent crawlers—streamlining the understanding of data assets. It’s being applied to tasks such as tagging, and extracting entities (people, places, and things) from unstructured data sources, which helps accelerate workflows and reduce manual overhead. I’m seeing automation tools that can design elements of the data architecture, assist with documentation, and manage refresh schedules.

On the analytics front, AI helps enhance user experiences by recommending visualizations, surfacing hidden insights in BI tools, and even generating natural language summaries of AI model results. AI is also increasingly important in model monitoring, helping to detect drift and maintain model performance over time.

Governance is another area where embedded AI is gaining traction. In addition to helping build trust in data via automated data quality checks, AI is helping organizations implement privacy policies, monitor data for anomalies, and streamline compliance updates. We'll discuss more about this in Part V.

Analytics and AI Investment Trends

Despite all the hype around AI, TDWI research indicates that only about half of the organizations we survey are currently using traditional AI. The other half remain earlier in their analytics journey, primarily focused on self-service analytics (such as data visualization) and dashboards. I've been asking the same survey question around priorities for the past seven years. I've seen about 30% of survey respondents adopting AI seven years ago, and this number is now up to about 45–50%.¹² In other words, it has been a long road for many organizations. This journey reflects a broader pattern we continue to observe: while the promise of AI is significant, many organizations are still laying the foundational groundwork needed to fully realize its potential.

That said, the momentum toward AI is growing. Generative AI, in particular, has rapidly risen to the top of the list of analytics priorities in TDWI surveys. In fact, in 2025, it ranked second only to self-service analytics in terms of organizational focus.¹³ The vast majority of organizations we survey are using it, albeit in the consumerized fashion described earlier. Traditional machine learning, which has been the core AI technique for over a decade, ranks behind generative AI in current prioritization. This suggests that many organizations may be skipping over more incremental advances in predictive analytics, possibly due to the perceived accessibility and transformative potential of generative AI tools. As a result, they may be missing out on the valuable insights and business learnings that more traditional machine learning applications can offer.

Importantly, when we ask organizations why they are modernizing their data infrastructures, the top motivation is to support AI initiatives. Whether they are building scalable data environments needed for training models or improving data quality and governance to ensure trustworthy AI inputs and outputs, organizations are aligning their modernization efforts with the goal

¹² TDWI, unpublished survey results, 2025.

¹³ TDWI, unpublished survey results, 2025.

of becoming AI-ready. This trend underscores an important insight: even if AI adoption is not yet widespread, the strategic imperative to prepare for AI is driving significant investment and reshaping data strategies across organizations.

Justifying Investments with Concrete Outcomes

To successfully justify AI investments, organizations must tie those initiatives to measurable business problems and outcomes. It's much easier to make a case for AI when it is clearly solving a specific challenge, such as improving retention, reducing downtime, or increasing sales conversion, than when it's seen as an experiment in search of a use. That is why bringing stakeholders together at the start is so important. They can help frame business priorities, identify the right use cases, and champion the solution once it's deployed. This collaborative approach creates organizational alignment and reduces resistance to AI adoption. Chapter 13 digs much deeper into tying metrics to AI value.

AI should not be a “build it and they will come” initiative. Instead, it should be grounded in the business—serving a strategic purpose, delivering real value, and solving clearly defined problems.

What's Next?

The next few chapters of Part I explore the importance of strong leadership and why it is essential for successful AI adoption. We'll look at the pitfalls of visionary strategies that fail without solid execution, and the critical need for alignment between business goals and technical expertise. We'll then map out the organizational journey toward AI maturity, introducing a model, used by TDWI, that captures the five key dimensions of readiness.

