Microsoft Azure

Generative AI Lessons Learned

Discover real-world use cases and outcomes with Azure AI Foundry across industries

Contents

Introduction

Building the next generation of intelligent apps

Generative AI isn't a trend—it's the future. AI adoption has surged from 50% to 72% in just one year, revealing a growing global interest. As industries rush to embrace this technology, the drive to build the next generation of intelligent apps has never been stronger.

But the path to success isn't always clear. Many organizations struggle to determine the best way to leverage generative AI effectively. While the excitement is high, real-world challenges remain. Companies need to address several key issues to build truly transformative AI applications.

Roadblocks to building impactful generative AI apps

- **Ensuring high-quality data:**
  Poorly prepared or uncurated data can result in biased or inaccurate AI outcomes, reducing the effectiveness of AI initiatives.
- **Customizing off-the-shelf AI solutions:**
  Many generative AI tools don't fully align with an organization's unique needs.
- **Filling expertise gaps:**
  Building and deploying AI solutions often requires specialized skills that may be in short supply.
- **Demonstrating clear value and ROI:**
  Many organizations struggle to calculate and communicate the return on investment for AI initiatives.
- **Addressing security and compliance concerns:**
  Handling sensitive data with AI introduces risks, requiring robust security protocols and compliance measures to safeguard data.
- **Integrating generative AI into existing systems:**
  Significant upgrades or system redesigns may be required to avoid disruptions and achieve maximum efficiency.

These roadblocks can make the process of developing and deploying AI applications more complex than anticipated, limiting the app's potential to deliver the desired outcomes. However, your team can significantly ease these burdens and accelerate innovation with the right tools and resources.

Azure AI Foundry: An AI playground for the business world

Azure AI Foundry is a centralized development hub for building safe, scalable generative AI applications. Developers can tap into the latest, cutting-edge models from the open-source community and strategic Microsoft partnerships, as well as Azure AI Services and pre-built templates.

Azure AI Foundry is designed to help organizations overcome the common challenges of adopting and implementing generative AI by offering a wide range of tools and models. Whether you're planning to use pre-built models or build custom solutions, your team gets guided support for model benchmarking, solution evaluation, prompt engineering, fine-tuning, content safety, continuous monitoring, and more.

- **Azure AI Foundry:** Design, customize, and manage generative AI applications with prebuilt and customizable models
- **Azure Machine Learning:** Use tools for data preparation, model training, and deployment to create an end-to-end machine learning lifecycle.
- **Azure AI Agent Service:** Create, deploy, and monitor AI agents that can automate complex business processes.

Azure AI Services

- **Azure OpenAI Service:** Build custom generative AI solutions with advanced models like GPT-4 and DALL-E.
- **Azure AI Search:** Implement enterprise-ready search and retrieval systems with advanced search technologies.
- **Azure AI Content Safety:** Enhance the safety of AI applications by detecting and blocking harmful content.
- **Azure AI Translator:** Translate text instantly across more than 100 languages using neural machine translation.
- **Azure AI Speech:** Convert speech to text, text to speech, and translate spoken audio in real-time.
- **Azure AI Vision:** Analyze images and videos to extract text, detect faces, and recognize objects.
- **Azure AI Language:** Utilize natural language processing to understand and analyze text for various applications.
- **Azure AI Document Intelligence:** Extract text, key-value pairs, and tables from documents using advanced machine learning.

Curious to see what's possible with generative AI? Discover how teams across industries are applying generative AI to boost agility, enhance decision-making, and deliver superior customer experiences.

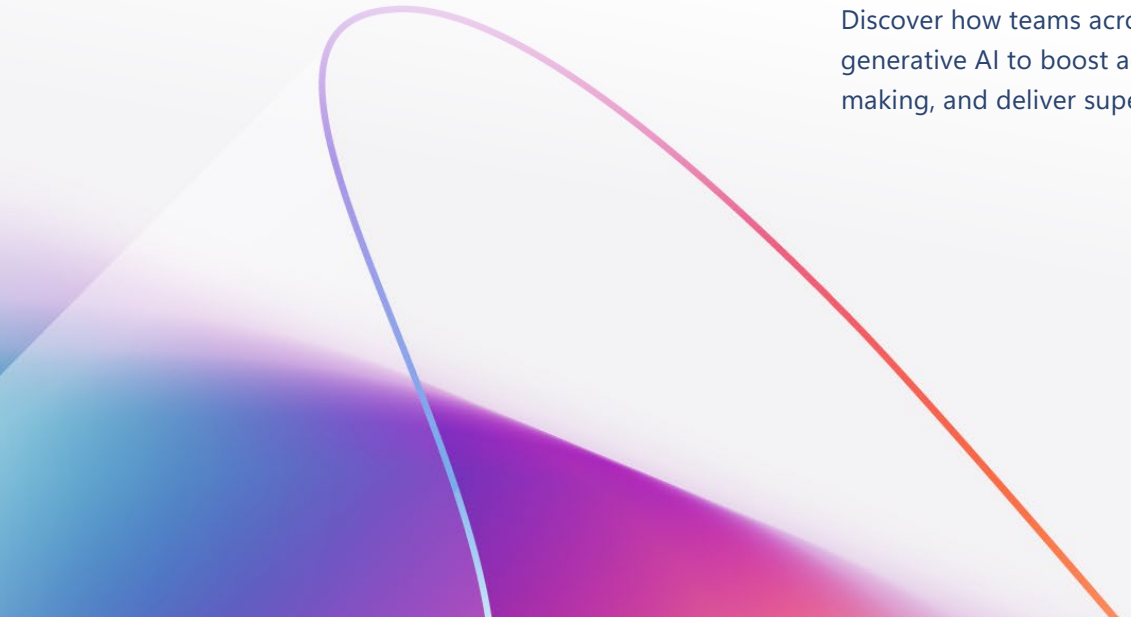

Chapter 1

Financial Services

Financial services companies face increasing regulatory pressures, the need to securely manage vast amounts of data, and a growing demand for personalized customer experiences. They must also navigate market volatility, improve operational efficiency, and safeguard against cyber threats, all while staying ahead of competitors in an increasingly digital world.

Generative AI can play a pivotal role in addressing these challenges by streamlining operations, enhancing decision-making, and enabling personalization at scale. Through advanced data analysis and automation, generative AI helps financial institutions uncover insights, automate routine tasks, and create more engaging customer experiences.

Generative AI use cases for financial services

Use case: Fraud detection

Sample use case: AI-powered fraud detection helps financial services organizations identify and prevent fraudulent activities. By analyzing incoming data and learning from historical patterns, AI enables operations teams and product leaders to detect anomalies such as irregular transaction amounts or unusual purchase locations. By helping improve the precision (accuracy of fraud identification) and recall (completeness of detection), AI-powered fraud detection has the potential to save organizations millions annually.

Use case: Wealth advisor AI assistant

A wealth advisor AI assistant offers financial professionals personalized support and data-driven insights. It creates tailored documents, proposals, and investment strategies for clients, enabling more meaningful client interactions. By reducing the time advisors spend on administrative tasks, the AI assistant allows them to focus on building client relationships and better understanding client needs. This improves client retention, customer satisfaction (CSAT), and investment performance.

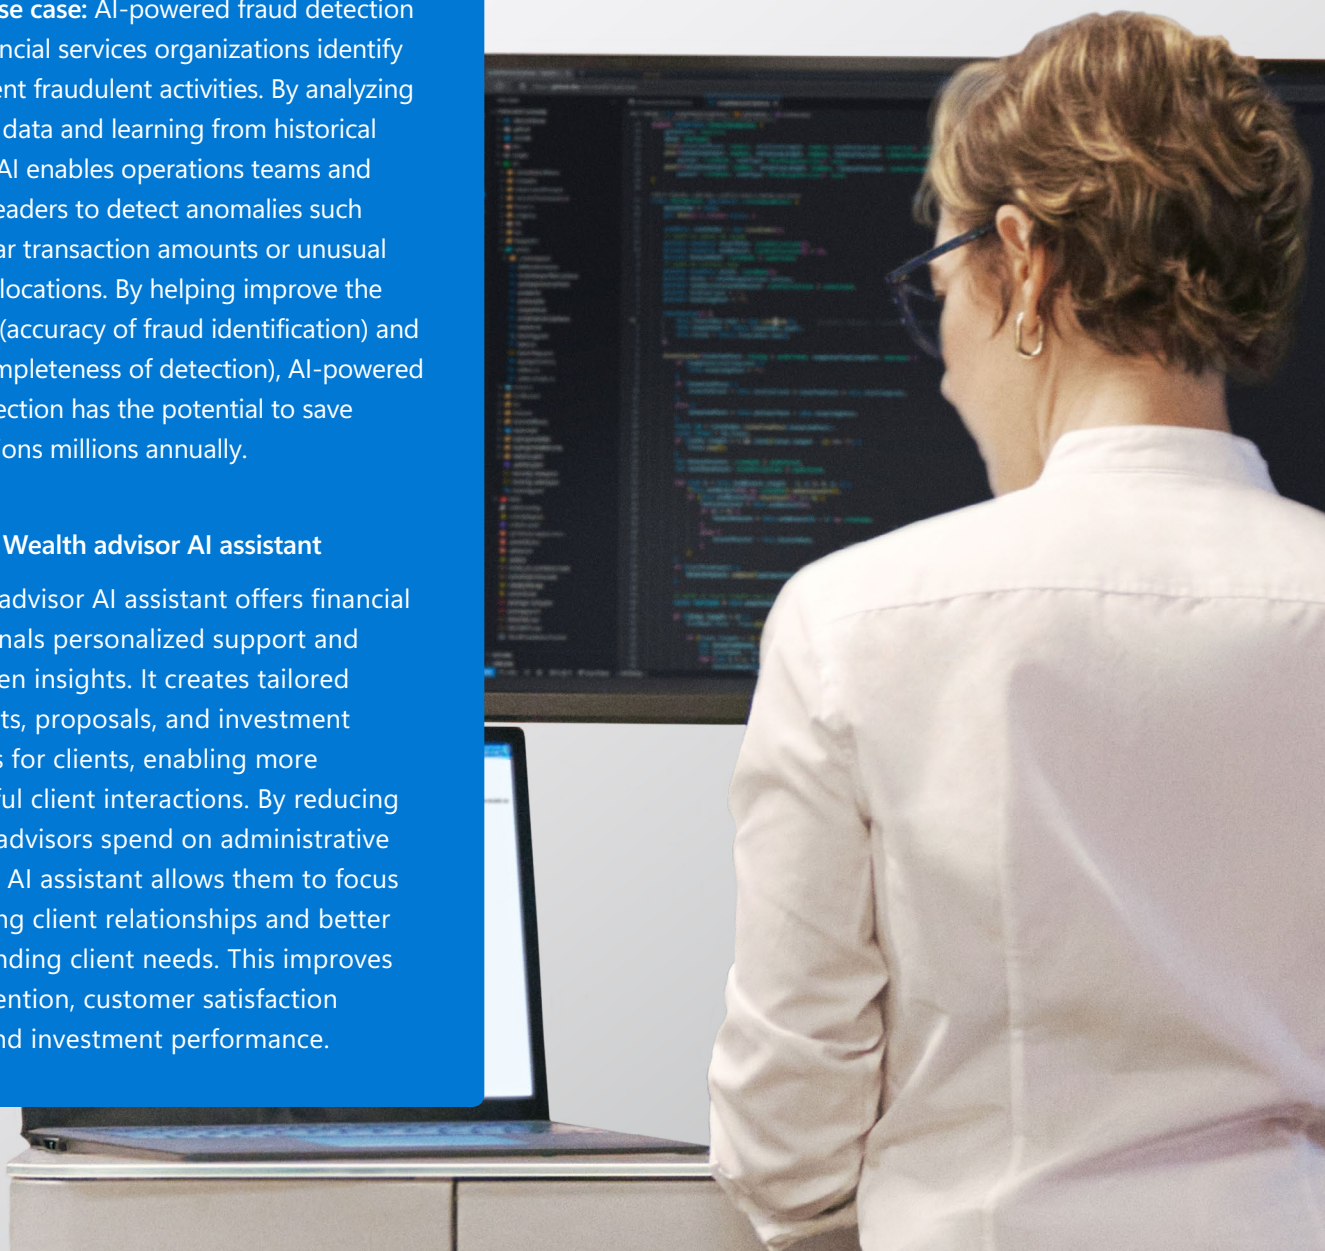

Zurich Insurance Group shapes the future of underwriting with Azure OpenAI Service

The challenge

With customers describing their needs in images, emails, and reports using their local language, Zurich Insurance Group recognized the challenge to translate unstructured data into a rules-based environment. Working to get information faster, Zurich wanted to transform how and how quickly the company serves its customers, making real-time access to data and analytics critical to its mission.

The solution

Zurich Insurance Group turned to tools such as Microsoft Azure OpenAI Service for its advanced language AI, text-to-speech models, and highly secure, enterprise-grade Azure features.

Zurich continuously innovates, and in deploying AI across its value chain, the organization is improving risk assessment and increasing personalization. Working closely with companies like Microsoft, Zurich is making a significant impact across its global, decentralized footprint, setting new standards within a more secure, trusted, scalable environment.

The outcome

Zurich Insurance Group is developing advanced AI applications that lead to more accurate and efficient risk management evaluations, accelerating the underwriting process, reducing turnaround times, and increasing customer satisfaction.

One of the key areas where Zurich is focusing its AI efforts is in underwriting. By adopting tools like Microsoft Azure OpenAI Service, Zurich is developing advanced AI applications that lead to more accurate and efficient evaluations. This not only improves the quality of risk assessments but also accelerates the underwriting process. The result is reduced turnaround times and increased customer satisfaction. Zurich continues to set new benchmarks for innovation, efficiency, and customer satisfaction. Bringing its ideas to fruition, Zurich has more than 200 AI tools currently deployed.

More to explore

Bradesco

Bradesco Bank improved manager productivity by integrating Microsoft Azure's generative AI into its virtual assistant, BIA. This allowed for faster and more accurate responses to internal regulations, **cutting response time from days to hours** and achieving an **82% resolution rate**.

Trusting Social

"With Azure Machine Learning and Azure OpenAI Service, we're able to build and train [our] generative AI agents with our own models and those available on the Azure AI platform to meet the specific needs of organizations."

—Ha Le, Director of Product Innovation at Trusting Social

Chapter 2

Healthcare

Healthcare organizations today are navigating a complex landscape marked by rising operational costs, increasing regulatory requirements, and growing demand for improved patient outcomes. They must manage vast volumes of sensitive data, ensure compliance with evolving privacy laws, and meet the expectations of patients who now seek personalized, seamless care experiences. In this demanding environment, AI offers transformative potential. By automating routine tasks, uncovering insights from complex datasets, and enhancing decision-making, AI can help healthcare organizations tackle their most pressing challenges and deliver better care efficiently and securely.

Generative AI use cases for healthcare

Sales representative training and coaching

Pharmaceutical companies need to train and coach their sales representatives on product knowledge, compliance, and sales techniques. An AI-powered training system automates personalized content delivery, tracks progress, and provides real-time feedback, transforming unstructured training data into easily accessible digital information. This enables faster, more consistent, and reliable training, improving sales performance, ensuring compliance, and better preparing representatives for healthcare professional engagement.

Research and development semantic search

Pharmaceutical companies need to manage large volumes of research and development documents. A semantic search system automates the indexing, categorization, and retrieval of documents based on their content, turning unstructured data into accessible, structured information for efficient analysis and storage. This system improves document accessibility, accelerating research, enhancing regulatory compliance, and ensuring critical information is readily available for decision-making. It also reduces the time and cost of manual document management.

NHS saves clinicians time by making data more discoverable

The challenge

Cambridgeshire and Peterborough NHS Foundation Trust needed to provide clinicians with fast, secure access to 23 years of patient data stored across multiple systems. With over three million clinical documents to migrate and a limited timeline, traditional migration methods were either too costly or time-consuming. Without a reliable way to search across these datasets, clinicians faced delays in accessing critical information, impacting patient care.

The solution

The Trust turned to Azure AI Search to consolidate and unlock insights from its vast data archives. The internal IT team developed a solution that indexed legacy data, structured clinical forms, and handwritten notes. Using advanced search capabilities like Optical Character Recognition (OCR) for handwritten PDFs, clinicians gained seamless access to comprehensive patient records via a secure web interface. The system is also integrated with Microsoft Entra ID for effortless and secure user authentication.

The outcome

- Clinicians can now search across millions of documents in seconds, improving response times and patient outcomes.
- Fast, accurate search functionality transformed how clinicians access patient histories.
- The scalable Azure platform allows the Trust to expand its use case to other systems, including finance and employee records.

More to explore

Discover the results other healthcare organizations have been making with generative AI solutions built with Azure AI Foundry tools and solutions.

Kry

"By transcribing doctors' notes, pre-filling referral letters, and performing administrative tasks more efficiently, AI allows Kry clinicians to focus on delivering better care, while ensuring patients can access the advice, care, and treatment they need most efficiently."

—Fredrik Nylander, Chief Technology Officer at Kry

Doctolib

One of the leading eHealth companies in France, Doctolib used Azure AI Foundry to build an AI-powered assistant that can **summarize medical consultations in 15 seconds**, allowing practitioners to spend twice as much time with patients.

Partner spotlight:
Meta

The Azure AI Model Catalog offers a diverse selection of AI models, featuring options from Mistral and Meta, giving Azure customers the flexibility to select the model that best fits their specific needs.

Azure + Meta Llama models

Meta is committed to advancing AI with open-weight models, which allow researchers and developers to access and fine-tune publicly available trained weights. In collaboration with Microsoft, Meta has integrated its state-of-the-art Llama 3 models, such as Meta-Llama-3-8B-Instruct and Meta-Llama-3-70B-Instruct, into the Azure AI Model Catalog, empowering developers to leverage these tools for customization and innovation.

Discover how integrating Meta's Llama models into Azure AI enables developers to

- Improve reasoning across multiple domains so models can dynamically adapt to real-time conditions.
- Enable autonomous machine intelligence (AMI), allowing models to understand cause and effect.
- Accelerate inference speeds and reduce memory usage.
- Use models for scenarios with limited computational resources, such as text summarization, classification, sentiment analysis, and translation.

Customer use case

UC Berkley uses Llama models on Azure to improve Retrieval Augmented Generation (RAG) performance, enabling a better way to create natural language interfaces with access to existing knowledge.

Chapter 3

Manufacturing

From supply chain disruptions and fluctuating raw material costs to increasing demands for sustainability and compliance with stringent safety regulations, manufacturers have no shortage of challenges. They also grapple with optimizing production efficiency, reducing downtime, and meeting the growing demand for customized products—all while staying ahead in a rapidly digitizing market. Labor shortages and the need to upskill workers for advanced technologies add further pressure.

Amid these challenges, generative AI offers powerful solutions. AI empowers manufacturers to enhance efficiency, reduce costs, and adapt to changing demands by enabling predictive maintenance, improving supply chain visibility, and optimizing production processes. It also supports innovation, allowing manufacturers to stay competitive in an evolving industry.

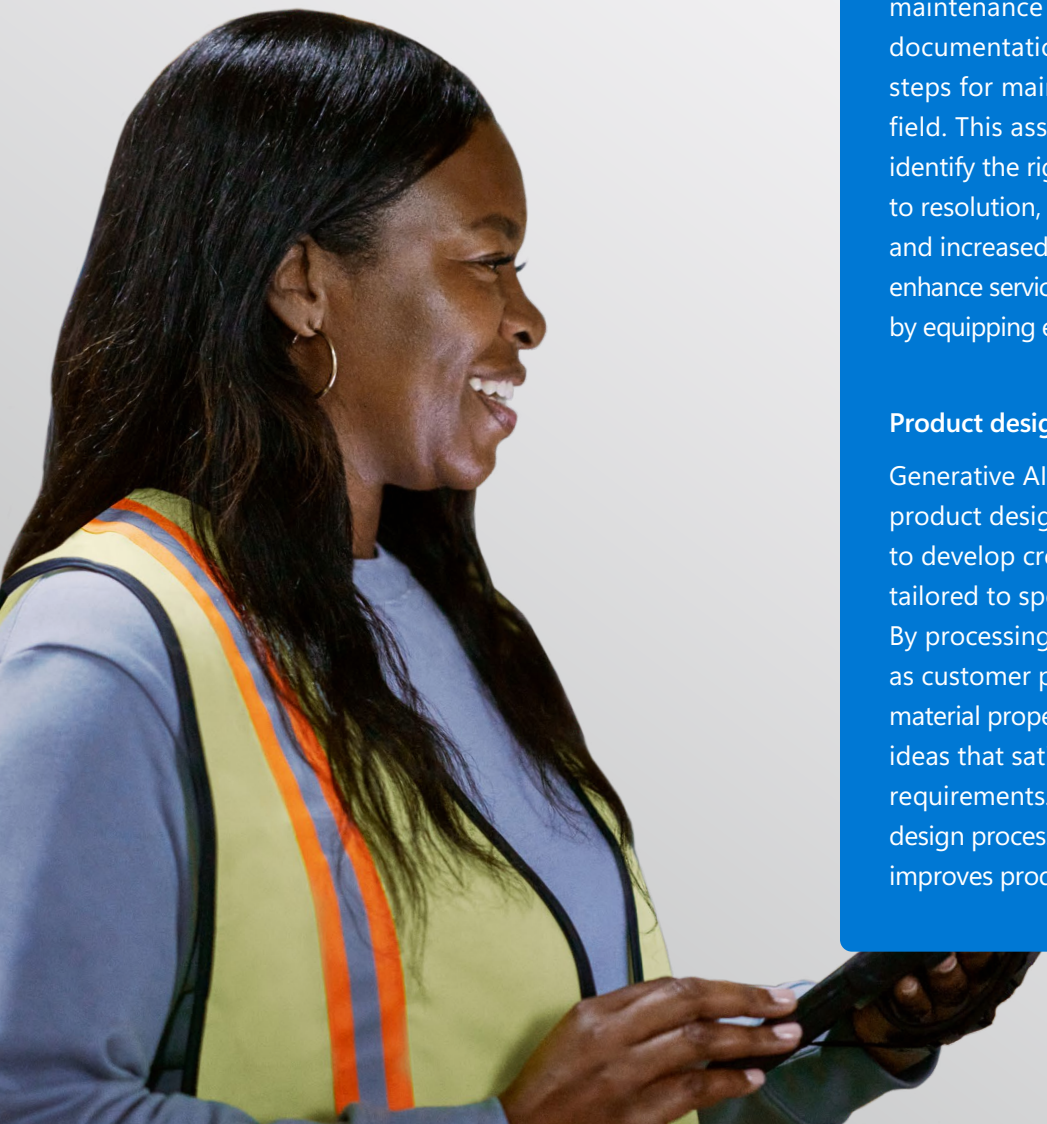

Generative AI use cases for manufacturers

AI field service assistant

The AI field service assistant supports field service engineers by combining historical maintenance and repair data with product documentation to suggest the next best steps for maintenance and repairs in the field. This assistant helps engineers quickly identify the right actions, enabling faster time to resolution, improved customer experience, and increased product uptime. Companies can enhance service quality and operational efficiency by equipping engineers with this AI tool.

Product design

Generative AI has the potential to transform product design by using advanced algorithms to develop creative and optimized designs tailored to specific needs and limitations. By processing large amounts of data, such as customer preferences, market trends, and material properties, it can propose unique design ideas that satisfy both functional and aesthetic requirements. This technology speeds up the design process, lowers development costs, and improves product performance and appeal.

Siemens enables real-time problem-solving for frontline workers and engineers

The challenge

Siemens Digital Industries Software recognized the need to improve communication across product lifecycle teams for its manufacturing customers. Field and shop floor workers often struggled to report issues due to language barriers, inconvenient interfaces, or fast-paced workflows. This lack of real-time feedback delayed issue resolution and stifled innovation.

The solution

Siemens collaborated with Microsoft to create an AI-powered app integrated with Microsoft Teams and its Teamcenter product lifecycle management (PLM) solution. Leveraging Azure AI Foundry natural language processing capabilities, the app allows workers to report problems using voice in their native language. The app translates and summarizes the input, routing it to the appropriate team members in the required language.

The outcome

- Workers can report issues in real-time, leading to faster problem resolution.
- Improved cross-functional collaboration by breaking down language and technical barriers.
- Increased inclusion and engagement of frontline workers in design and manufacturing processes.
- Enhanced workflows and innovation cycles for Siemens customers.

More to explore

See how other manufacturers have been innovating with generative AI solutions built with Azure AI Foundry tools and solutions.

Schneider Electric

Schneider is enhancing productivity across the organization by extending Azure OpenAI Service, with applications like its Knowledge Bot, which helps customer service agents quickly find the information they need to answer customer queries. Financial analysts at Schneider also use Finance Advisor to support more informed and optimal decision-making.

"Generative AI solutions like Azure OpenAI are delivering what I dreamed of when I started my career: **seamless cooperative interaction between human and machine** to make the best possible decisions."

—Francois Bonnard, AI Partnerships Director at Schneider Electric

Chapter 4

Retail

Retail companies face unique challenges in today's fast-paced and competitive environment. These include adapting to shifting consumer expectations, managing complex supply chains, and delivering personalized experiences across multiple channels. Retailers must also contend with increasing pressure to embrace sustainability, navigate fluctuating market trends, and ensure secure, seamless payment systems.

Generative AI has the potential to address these issues by enabling smarter inventory management, optimizing pricing strategies, and delivering hyper-personalized customer experiences. AI-powered tools can also improve demand forecasting, streamline supply chain operations, and enhance fraud detection, helping retailers remain competitive in a dynamic market.

Generative AI use cases for Retail

Hyper-personalization and product discovery

If customers feel it takes too long to find the product they want, they may become frustrated and abandon their search, ultimately leading to a lost sale. AI-powered product discovery helps customers find exactly what they're looking for through semantic search. By analyzing customer behaviors and preferences, AI tailors shopping experiences, offering personalized product recommendations and promotions, leading to more sales and higher customer satisfaction and loyalty. With more relevant recommendations, retailers can drive engagement with customers more likely to return, improving retention and engagement over time.

Retail associate AI assistant

AI assistants support retail associates by providing real-time information about products, inventory, and customer preferences. By helping with tasks such as restocking and locating items, these assistants enable associates to offer better service and make informed decisions, improving overall store operations. AI assistance also supports a more positive work environment, reducing turnover and increasing employee retention.

ASOS gives shoppers tailored experiences that fit unique tastes

The challenge

ASOS, a leading UK-based online fashion retailer, sought to enhance its customer experience by delivering hyper-personalized shopping options. With a diverse catalog spanning nearly 900 brand partners and its own labels, the company wanted a solution to help customers quickly discover relevant trends and styles. Existing search and browse tools were effective but lacked the conversational, interactive touch that tech-savvy customers expect. ASOS also needed a way to integrate real-time fashion trends into the shopping journey while maintaining its distinctive brand voice and high ethical standards.

The solution

ASOS turned to Azure AI Foundry, Azure OpenAI Service, and Azure AI prompt flow to develop a conversational AI-powered shopping experience. This innovative solution enables customers to engage naturally with ASOS on its website and app to find products tailored to their preferences. Using ChatGPT language models, the experience curates fashion recommendations based on individual tastes and emerging trends.

ASOS built a prototype in just weeks, using prompt flow to streamline development, ensure secure prompt engineering, and align the AI's tone with the ASOS brand. The solution integrates seamlessly with internal systems, pulls trend data from various sources, and adapts dynamically to evolving customer needs. ASOS also prioritized responsible AI practices to ensure the experience aligns with its ethical standards, reducing bias and safeguarding customer interactions.

The outcome

- Enabled customers to discover curated fashion trends in a natural, conversational way.
- Created a scalable, AI-driven solution that aligns with ASOS's innovative brand image.
- Strengthened ethical standards by incorporating Azure AI Content Safety for responsible interactions.
- Accelerated development cycles, reducing the time to prototype from months to weeks.

More to explore

Retail companies around the world have been implementing generative AI to help deliver better shopping experiences.

Iceland

Supermarket chain Iceland built a new app called Genie which leverages the natural language capabilities of Azure OpenAI so colleagues can interact with business knowledge. Now, staff can ask natural language questions of Genie to understand what they should do.

"It's so much faster than a traditional search. Instead of being presented with links to 15 different PDFs and then having to scan them all to see if one contains the answer you need, you are **presented with the answer immediately**,"

—Lee Boswell, Consumables and Process Manager at Iceland

Partner spotlight:
Mistral AI

Azure + Mistral Large models

Mistral AI focuses on open-source innovation and exceptional model performance, offering advanced language models like Mistral-Large and Mixtral-8x7B. These models deliver high performance, scalability, and flexibility for tasks such as conversational AI and coding. Available through Azure AI Foundry and Azure Machine Learning Model Catalog, they can be deployed as serverless APIs, via Azure AI Studio, or on virtual machines, providing efficient, low-latency solutions for diverse use cases.

Learn how using Mistral Large models helps developers to

- Fine-tune models on specific data to enhance performance for personalized marketing generation, summarization, and customer support.
- Enable multilingual applications.
- Optimize coding tasks for specific industries like video games, healthcare, and legal.
- Ensure models are compliant with industry standards.
- Use advanced functionalities like RAG, synthetic text generation, and multi-lingual virtual assistants.

Customer use case

A global leader in insurance and asset management, AXA uses Mistral AI's models to speed up response times in customer interactions, enable more accurate risk assessments, and increase customer engagement.

Chapter 5

Other industries

The possibilities for generative AI are endless, and these are just a few of the many ways it's transforming industries. Discover how organizations in other sectors are putting generative AI to work.

Professional Services

One of Europe's leading architecture and engineering firms, Sweco, wanted to automate documentation creation and analysis and enhance search capabilities to help consultants get the needed information.

"Through simple coding and logic creation through prompt flow, [Azure AI Foundry (formerly Azure AI Studio)] made it easy for us to meet the specific needs of our colleagues... Within two weeks of the initial POC, SwecoGPT was complete, and we deployed it. Users started using it the same day and loved it."

—Shah Muhammad, Head of AI Innovation, Sweco

Telecommunications

Global telecommunications company Vodafone used Azure AI Foundry and Azure OpenAI to build two generative AI solutions—SuperTOBi and SuperAgent—to handle complex customer inquiries and free customer care agents from repetitive tasks. Initial tests of SuperTOBi have shown a 50% improvement in first-time resolution with tasks like billing inquiries.

Government

The South Australia Department for Education launched the EdChat chatbot, using Azure AI Content Safety to protect students from harmful content. Initially, 20% of students used EdChat during the trial, but usage grew as teachers and students recognized its potential for critical and creative thinking. The built-in safety features, which block inappropriate queries and filter harmful responses, allowed teachers to focus more on the tool's educational benefits than on content oversight.

Driving innovation with collaboration

To stay at the forefront of AI innovation, Azure collaborates with other leading tech companies, combining expertise, resources, and cutting-edge technologies to accelerate the development and deployment of AI solutions. These partnerships enable Azure to integrate advanced AI tools, frameworks, and services that enhance machine learning capabilities, improve data analytics, and foster the creation of intelligent applications.

Azure + Freshworks: Building a better way to sell, market, and develop

Freshworks enhanced its Freddy AI Copilot by integrating Azure OpenAI Service and Azure AI Services, including the OpenAI Completions API, GPT-3.5-turbo, and GPT-4 models, to implement features like message expansion, tone enhancement, summarization, rephrasing, and more.

Launch your generative AI journey with Azure AI Foundry

As these customer stories show, generative AI has the power to drive immense transformations for organizations across industries. However, building these solutions often comes with challenges, such as selecting the right models, managing costly infrastructure, and handling complex data. Teams must also ensure their AI applications are safe, reliable, and compliant with data governance.

Azure AI Foundry helps overcome these obstacles by offering a centralized hub where developers can access cutting-edge models, use pre-built templates, and receive support for model selection, benchmarking, fine-tuning, and content safety. With Azure AI Foundry, your organization can streamline development, reduce infrastructure costs, and confidently build scalable, innovative, and secure AI applications that remain adaptable in the rapidly evolving AI landscape.

Learn more about how Azure AI Foundry puts you on the fast track to building or customizing your own generative AI applications.

Explore Azure AI Foundry →