

Industry-Wise Use Cases of Generative AI (GenAI)

Part 01 — Week 01 Research Document

Prepared by: Shivya

Organization: MacroEdtech

Program: GenAI Research Internship — Summer 2026

Submission Date: June 12, 2026

5 Industries

47 Use Cases

10+ Patterns

25 References

Table of Contents

Table of Contents	2
About This Document	4
1. Introduction to Generative AI	5
1.1 What is Generative AI?.....	5
1.2 How Does GenAI Work? (Simple Explanation).....	5
1.3 Common Types of GenAI Models.....	5
1.4 Why is GenAI Important?	6
1.5 Industries Covered in This Report	6
2. Healthcare & Life Sciences	7
2.1 Use Case Overview Table	7
2.2 What I Found Most Interesting.....	8
Clinical Documentation — Solving Burnout.....	8
Drug Discovery — AI Changing Science.....	8
2.3 My Key Observations from Healthcare	8
3. Banking & Financial Services.....	10
3.1 Use Case Overview Table	10
3.2 Fraud Detection — Understanding How It Works	11
3.3 Important Regulations Affecting AI in Banking.....	11
3.4 My Key Observations from Banking.....	12
4. Manufacturing & Industry 4.0	13
4.1 Use Case Overview Table	13
4.2 Digital Twins — What They Are and Why They Matter	14
4.3 Industry 4.0 — How the Layers Work Together	14
4.4 My Key Observations from Manufacturing.....	14
5. Retail & E-Commerce	16
5.1 Use Case Overview Table.....	16
5.2 Why Recommendations Are Such a Big Deal.....	17
5.3 How These Use Cases Compare in ROI	17
5.4 My Key Observations from Retail	18
6. Information Technology & Software	19
6.1 Use Case Overview Table	19
6.2 What GitHub Copilot Is Teaching Us About AI Productivity	20
6.3 A Note on Security — AI Can Create New Risks Too.....	20
6.4 My Key Observations from IT & Software.....	20
7. Patterns I Noticed Across All Industries	22
7.1 Pattern 1: Automating Documents and Reports.....	22
7.2 Pattern 2: Making Complex Systems Easier to Use.....	22

7.3 Pattern 3: Predicting Problems Before They Happen	22
7.4 Pattern 4: Personalization at Large Scale.....	22
7.5 Pattern 5: Helping Organizations Stay Compliant.....	23
7.6 Challenges I Found — Why GenAI Is Not Perfect	23
8. Future Trends — What I Think Is Coming Next.....	24
8.1 Agentic AI — AI That Takes Actions on Its Own.....	24
8.2 Multimodal AI — AI That Understands Everything at Once.....	24
8.3 Domain-Specific AI Models	24
8.4 Better Collaboration Between Humans and AI.....	24
8.5 Real-Time AI Everywhere.....	25
9. Conclusion and Personal Learnings.....	26
9.1 My Top 5 Personal Learnings.....	26
9.2 What I Would Like to Research Further	26
Appendix A: GenAI Tools I Explored During Research	28
Appendix B: Key Terms and Definitions (Glossary).....	29
10. References	31
10.1 Academic Papers	31
10.2 Industry Reports	31
10.3 Company and Product Documentation	32
10.4 Regulatory and Governance Sources.....	32

About This Document

This document is my Week 01 research submission for the MacroEdtech GenAI Research Internship, Summer 2026. The task assigned to me was to study real-world use cases of Generative AI (GenAI) across different industries and prepare a detailed research document based on my understanding.

Over the course of four days (June 9–12, 2026), I studied the reference material shared in our Google Drive, explored additional resources, and documented my findings in this report. The goal was to understand how GenAI is actually being used in industries like Healthcare, Banking, Manufacturing, Retail, and IT — not just in theory, but in real companies solving real problems.

This report covers 47 use cases across 5 industries. For each industry, I have explained what GenAI is doing, what impact it is creating, and which companies are using it. I have also shared my observations about patterns I noticed across industries, challenges that companies face, and what I think the future looks like.

What I Was Assigned to Do (Week 01 Task Summary)

- Study the Part 01 — GenAI Use Cases reference material shared in Google Drive
- Research and document use cases of GenAI across 5 industries
- Understand how GenAI is solving real business problems
- Identify patterns, challenges, and future trends
- Submit as a detailed research document in PDF or PPT format

1. Introduction to Generative AI

Before I started researching industry-specific use cases, I wanted to make sure I clearly understood what Generative AI actually is. This section covers the basics that I learned during my study.

1.1 What is Generative AI?

Generative AI (GenAI) is a type of Artificial Intelligence that can create new content on its own — such as text, images, code, audio, and data — by learning patterns from large amounts of existing data. What makes it different from older AI systems is that instead of just predicting a label or category (like 'is this email spam or not'), GenAI generates something completely new.

For example, if you ask ChatGPT to write a summary of a medical report, it does not copy anything — it understands the content and generates a new, readable summary. That is the core idea behind Generative AI.

1.2 How Does GenAI Work? (Simple Explanation)

I found it helpful to understand the basic process of how these models are built and trained:

1. **Pre-training:** The AI model is first trained on a huge amount of text (books, websites, articles, code) so it learns how language and information work.
2. **Fine-tuning:** After basic training, the model is further trained on specific types of data (for example, medical records) so it becomes better at that particular field.
3. **RLHF (Reinforcement Learning from Human Feedback):** Human reviewers rate the AI's answers, and the model learns to give better, more helpful responses based on this feedback.
4. **Prompt Engineering:** Users give the AI clear instructions (called prompts) to get useful outputs. The quality of the prompt directly affects the quality of the output.

The technology behind most modern GenAI systems is called the Transformer architecture, which was introduced in a famous 2017 research paper called 'Attention Is All You Need' by Vaswani et al. This architecture allows AI to understand context across long pieces of text, which is why it performs so well.

1.3 Common Types of GenAI Models

Type of Model	What it Can Do / Where it is Used
Large Language Models (LLMs)	Text generation, summarization, Q&A, document writing, code generation
Vision-Language Models	Medical image analysis, product search by photo, document scanning
Code Generation Models	Writing software code, fixing bugs, generating test cases
Diffusion Models	Creating realistic images, product designs, synthetic data
Multimodal Models	Working with text + images + audio together in one system

Type of Model	What it Can Do / Where it is Used
Domain-Specific LLMs	Specialized models for medicine (BioMedLM), finance (BloombergGPT), etc.

1.4 Why is GenAI Important?

From what I studied, here are the main reasons GenAI is becoming so important across industries:

- It automates tasks that require reading, writing, and understanding language — which were previously very hard to automate
- It helps organizations process huge amounts of unstructured data (documents, images, audio) much faster than humans can
- It allows companies to offer personalized experiences to millions of customers at once, at very low cost
- It speeds up research and development in fields like medicine, science, and engineering
- It makes complex systems easier to use by allowing people to interact with them in plain language

1.5 Industries Covered in This Report

- Healthcare & Life Sciences — How GenAI is helping doctors, researchers, and patients
- Banking & Financial Services — How GenAI is improving security, customer service, and decision-making
- Manufacturing & Industry 4.0 — How GenAI is making factories smarter and more efficient
- Retail & E-Commerce — How GenAI is improving shopping experiences and supply chains
- Information Technology & Software — How GenAI is changing how software is built and maintained

2. Healthcare & Life Sciences

Healthcare was one of the first industries to start seriously using Generative AI, and it is also one of the areas where the impact is most significant. The reason is simple — healthcare involves a massive amount of written data (patient records, medical reports, research papers, prescriptions) that is very time-consuming for doctors and nurses to manage manually.

According to McKinsey (2024), Generative AI could create over \$1 trillion in value for the global healthcare industry. The AI in healthcare market is expected to cross \$188 billion by 2030. I was surprised to learn that doctors in the US spend an average of 4.5 hours every day just on documentation — more time than they actually spend with patients. GenAI is directly addressing this problem.

2.1 Use Case Overview Table

Use Case	What GenAI Does	Business Impact	Real-World Example
Clinical Documentation Automation	Listens to doctor-patient conversations and automatically writes clinical notes, discharge summaries, and reports.	Doctors save 50–70% of their documentation time, reducing burnout and spending more time with patients.	<i>Nuance DAX (Microsoft)</i>
Medical Imaging Report Generation	Analyzes X-ray, CT scan, and MRI images to automatically generate structured radiology reports.	Speeds up diagnosis and catches problems more consistently than manual review.	<i>Google Health, Aidoc</i>
AI-Assisted Drug Discovery	Studies how molecules interact to predict good drug candidates, reducing the need for physical lab testing.	Drug discovery time reduced from years to months; major cost savings in R&D.	<i>Insilico Medicine, Exscientia</i>
Synthetic Patient Data Generation	Creates realistic but fake patient data that researchers can use for AI training without privacy risks.	Enables AI research without violating patient privacy rules like HIPAA.	<i>Syntegra, MDClone</i>
Personalized Treatment Recommendations	Studies a patient's medical history, genetics, and current condition to suggest the best treatment plan.	Helps doctors choose more effective treatments and reduces trial-and-error.	<i>IBM Watson Health</i>
Biomedical Research Summarization	Reads thousands of research papers and produces short summaries of key findings for researchers.	What used to take weeks of reading can now be done in hours.	<i>Elicit, Semantic Scholar</i>
Healthcare Virtual Assistants	AI chatbots available 24/7 to answer patient questions, book	Reduces the pressure on hospital call centers; patients get faster responses.	<i>Babylon Health, Ada Health</i>

Use Case	What GenAI Does	Business Impact	Real-World Example
	appointments, and do basic health checks.		
Medical Coding Automation	Converts doctor's notes into standardized billing codes (ICD-10, CPT) automatically.	Reduces billing errors and speeds up hospital revenue cycles.	<i>Optum, 3M Health</i>
Genomic Data Interpretation	Reads genomic (DNA) data to identify health risks and recommend personalized care.	Makes precision medicine possible at scale.	<i>Tempus, Foundation Medicine</i>
Clinical Trial Optimization	Helps design clinical trials better, selects the right patients faster, and predicts outcomes.	Trials complete faster and cost less, improving drug approval rates.	<i>Medidata, Unlearn.AI</i>

2.2 What I Found Most Interesting

Clinical Documentation — Solving Burnout

The use case that stood out to me the most in healthcare was Clinical Documentation Automation. I was genuinely surprised to learn that doctors spend more time writing notes than talking to patients. Tools like Nuance DAX (built by Microsoft) listen to conversations between doctors and patients, and automatically produce structured medical notes without the doctor having to type anything. Doctors using this tool have reported that they can now focus fully on their patients during consultations, which even improves patient satisfaction.

Drug Discovery — AI Changing Science

Another use case I found fascinating was AI-Assisted Drug Discovery. Traditionally, developing a new drug takes 10–15 years and costs over \$2 billion on average. Most candidate drugs fail in clinical trials. GenAI is helping by predicting which molecules are likely to work as drugs before any physical experiments are done. A company called Insilico Medicine made history in 2023 by being the first company to take a fully AI-designed drug into Phase II clinical trials. That shows just how real this technology is becoming.

2.3 My Key Observations from Healthcare

Healthcare — What I Observed and Learned

- Documentation automation is growing fastest because it solves the very real and well-known problem of physician burnout
- AI drug discovery could cut development costs from \$2.6 billion to under \$1 billion per drug — this is a huge change for the pharmaceutical industry
- Synthetic data generation solves a unique problem in healthcare — you need patient data to train AI, but privacy rules prevent sharing it. Synthetic data is a clever solution

- Virtual health assistants are especially valuable in countries and regions where there are not enough doctors to serve the population
- One concern I noticed: AI models in healthcare can be biased if the training data doesn't include diverse patient groups — this is something researchers are still working to fix

3. Banking & Financial Services

Banking and finance is another industry that has been quick to adopt Generative AI. The reason is clear — banks handle enormous amounts of data every day (millions of transactions, thousands of customer interactions, huge volumes of regulatory documents) and even small improvements in speed or accuracy can translate into massive savings or revenue gains.

I was surprised to learn that global financial fraud losses exceed \$5 trillion every year. That alone explains why banks are so motivated to use AI for security. The AI in financial services market is projected to reach \$130 billion by 2030.

3.1 Use Case Overview Table

Use Case	What GenAI Does	Business Impact	Real-World Example
Fraud Detection via Anomaly Patterns	Monitors all transactions in real time and flags unusual activity that doesn't match a customer's normal behavior.	Fraud detection speed improved from hours to milliseconds; false alarms reduced by up to 50%.	<i>Mastercard Decision Intelligence</i>
Automated Financial Report Generation	Automatically creates financial reports, earnings summaries, and regulatory filings from raw financial data.	Saves hundreds of analyst hours per quarter; improves accuracy and consistency of reports.	<i>Bloomberg GPT, JPMorgan COiN</i>
AI-Powered Investment Research Assistants	Reads market data, news, and company filings to help analysts generate investment reports much faster.	Analysts can now cover 5x more companies with the same amount of time.	<i>Goldman Sachs Marcus, Morgan Stanley AI</i>
Regulatory Compliance Document Analysis	Reviews large volumes of regulatory documents to check if the bank is following all the rules.	Compliance review time reduced by 80%; fewer costly regulatory penalties.	<i>KPMG, Deloitte AI tools</i>
Intelligent Risk Assessment Models	Looks at a customer's credit history, market conditions, and other factors to assess lending risk.	More accurate loan decisions; lower default rates.	<i>ZestAI, Upstart</i>
Loan Underwriting Automation	Reviews loan applications, checks documents, and recommends approval or rejection automatically.	Loan approval time reduced from days to minutes.	<i>Blend, Ocrolus</i>
Financial Chatbot Advisors	AI chatbots answer customer queries, give budgeting advice, and recommend suitable financial products.	24/7 customer support without needing a human agent for every interaction.	<i>Bank of America Erica, Cleo</i>

Use Case	What GenAI Does	Business Impact	Real-World Example
KYC & AML Document Verification	Checks customer identity documents and screens transactions against fraud and money-laundering watchlists.	Customer onboarding time reduced from days to minutes; better compliance.	<i>Onfido, Sardine</i>
Market Sentiment Analysis	Reads news, social media, and earnings calls to understand how the market is feeling about a stock or sector.	Gives traders and analysts insights that are not available from traditional data sources.	<i>Two Sigma, Palantir</i>
Algorithmic Trading Support Systems	Generates and tests trading strategies using historical market data and pattern recognition.	Faster, more precise trade execution; strategies that adapt in real time to market changes.	<i>Citadel, Renaissance Technologies</i>

3.2 Fraud Detection — Understanding How It Works

I spent extra time trying to understand how GenAI-based fraud detection works differently from older systems. Traditional fraud detection used fixed rules — for example, 'flag any transaction above \$10,000' or 'block overseas transactions for this account.' The problem is that fraudsters learn these rules and work around them.

GenAI-based systems work differently. They learn the normal spending behavior of each individual customer — what they usually buy, when, where, and how much. If a transaction deviates significantly from that personal pattern, it gets flagged. This means the system can catch new types of fraud that were never seen before, because it is not looking for a specific pattern — it is detecting anything unusual for that customer.

Mastercard's Decision Intelligence platform uses this approach and processes over 100 billion transactions per year. It has improved fraud detection by 20% while reducing false alarms (which frustrate innocent customers) by 85%.

3.3 Important Regulations Affecting AI in Banking

One thing I learned that I didn't know before is how much regulation affects how banks can use AI. Banks can't just deploy any AI system they want — they have to follow strict rules. Here are the main ones I came across:

Regulation / Framework	What It Means for AI in Banking
EU AI Act (2024)	AI used for credit scoring and fraud detection must be explainable and have human oversight — the AI cannot just give a decision without justification
GDPR & CCPA	Banks must be careful about using personal customer data to train AI models; customers have the right to know why an AI made a decision about them

Regulation / Framework	What It Means for AI in Banking
Basel IV	Banks must prove that their AI risk models are accurate and traceable — they can't use a 'black box' model for capital calculations
SR 11-7 (US Federal Reserve)	All AI models used in banking must be validated, documented, and monitored on an ongoing basis

3.4 My Key Observations from Banking

Banking — What I Observed and Learned
<ul style="list-style-type: none">• Fraud detection is the most urgently needed use case — the sheer scale of global fraud (\$5 trillion annually) makes AI a necessity, not a luxury• I was impressed to learn that Bank of America's AI assistant 'Erica' has had over 1.5 billion customer interactions — showing that people are comfortable using AI for everyday banking• AI-generated financial reports raise an interesting ethical question: if a report is written by AI, who is responsible if there is an error? This is still being worked out• The explainability problem is real — many GenAI models cannot easily explain why they made a decision, which directly conflicts with banking regulations requiring transparency• Smaller fintech companies using AI-native systems are creating competitive pressure on traditional banks to adopt AI faster

4. Manufacturing & Industry 4.0

Manufacturing is one of the industries where I feel GenAI's impact is most tangible. In a factory, every hour of unexpected machine downtime can cost tens of thousands of dollars. Every defective product that slips through quality control has a cost. GenAI is being used to prevent these problems before they happen.

The concept of Industry 4.0 refers to the 'fourth industrial revolution' — where physical machines are connected to digital systems through the Internet of Things (IoT), and AI is used to analyze the data those machines produce. GenAI adds a new layer to this by generating insights, reports, and recommendations automatically from that data.

4.1 Use Case Overview Table

Use Case	What GenAI Does	Business Impact	Real-World Example
Predictive Maintenance Report Generation	Reads sensor data from machines (temperature, vibration, pressure) to predict failures before they happen.	Unplanned machine downtime reduced by up to 45%; machines last longer.	<i>Siemens MindSphere, GE Predix</i>
AI-Generated SOPs & Maintenance Manuals	Automatically creates and updates operating instructions and maintenance guides from machine logs and data.	Workers always have up-to-date guidance; less risk of accidents or mistakes.	<i>PTC, Rockwell Automation</i>
Smart Factory Digital Twins	Creates a virtual copy of a factory that mirrors the real factory in real time, used for testing and optimization.	New changes can be tested virtually before being applied to the real factory — saving time and money.	<i>NVIDIA Omniverse, Siemens</i>
Production Anomaly Detection	AI cameras and sensors watch the production line and alert workers when something unusual is detected.	Defects caught in milliseconds before an entire batch is ruined.	<i>Cognex, Landing AI</i>
Supply Chain Demand Forecasting	Predicts how much of each product will be needed based on historical sales, events, and market trends.	Less overproduction and waste; fewer stockouts.	<i>o9 Solutions, Blue Yonder</i>
Automated Engineering Documentation	Automatically generates technical drawings, bills of materials, and design documents from product specifications.	New products can be launched faster; fewer documentation errors.	<i>PTC Creo, Autodesk</i>
Root Cause Analysis Assistance	Analyzes production data to figure out why a quality problem occurred and suggests how to fix it.	Problem resolution time reduced from days to hours.	<i>Seeq, Sight Machine</i>

Use Case	What GenAI Does	Business Impact	Real-World Example
Inventory Optimization	Continuously adjusts stock levels across warehouses based on actual demand and delivery times.	Reduces warehousing costs by 20–30% while keeping products available.	<i>Llamasoft, Coupa</i>
Manufacturing Process Simulation	Runs computer simulations of entire manufacturing processes to find inefficiencies and test improvements.	Improvements can be safely tested without disrupting real production.	<i>AnyLogic, Arena Simulation</i>

4.2 Digital Twins — What They Are and Why They Matter

The concept of Digital Twins was new to me before I started this research, and I found it genuinely fascinating. A digital twin is basically a live, virtual copy of a real physical thing — it could be a single machine, a production line, or an entire factory. Sensors on the physical object constantly send data to the digital copy, keeping them synchronized.

The power of this is that you can run experiments on the digital twin without any risk to the real factory. For example, BMW uses digital twins to plan entire factory layouts before any physical construction happens. They can test different configurations virtually, find the most efficient one, and only then build it. This has reduced their factory reconfiguration time by 30% and virtually eliminated costly layout mistakes.

NVIDIA's Omniverse platform is one of the most advanced tools for creating industrial digital twins. Companies like Foxconn and BMW are already using it in production.

4.3 Industry 4.0 — How the Layers Work Together

Layer	Technologies Used	How GenAI Connects
Connectivity	IoT sensors, 5G networks	GenAI reads sensor data to detect anomalies
Data & Analytics	Data lakes, real-time streaming	GenAI finds patterns and generates predictions
AI & Intelligence	Computer vision, machine learning	Core GenAI models generate reports and decisions
Visualization	Digital twins, dashboards	GenAI generates natural language explanations of results
Automation	Robots, automated vehicles	GenAI directs and coordinates automated systems

4.4 My Key Observations from Manufacturing

Manufacturing — What I Observed and Learned

- Predictive maintenance is probably the use case with the clearest, most measurable ROI — if a machine costs \$50,000/hour in downtime and AI prevents even 2 failures per year, that's already \$100,000 saved
- Digital twins impressed me the most — the idea of testing a factory design virtually before building it feels like something out of science fiction, but it is already happening
- AI-generated documentation is solving a real workforce problem: many experienced manufacturing workers are retiring, and GenAI can capture their knowledge in documents before they leave
- Supply chain forecasting became critical after the COVID-19 pandemic exposed how fragile just-in-time manufacturing really is
- One challenge unique to manufacturing: AI systems need to be extremely reliable. A wrong prediction in a hospital is dangerous; a wrong prediction on a factory floor could injure workers or destroy equipment

5. Retail & E-Commerce

Retail was probably the industry I found easiest to relate to because I have personally experienced many of these GenAI applications as a customer. When Amazon recommends a product you were just thinking about, when a chatbot on a shopping website helps you find the right size, or when a fashion app lets you virtually try on clothes — that is all GenAI at work.

The global e-commerce market is now over \$6 trillion per year. With so many products and customers, retailers need AI to manage the scale. According to Salesforce (2024), about 17% of all e-commerce revenue is already influenced by AI recommendations.

5.1 Use Case Overview Table

Use Case	What GenAI Does	Business Impact	Real-World Example
Personalized Product Recommendation Engines	Studies what you've browsed, what you've bought, and what similar customers liked to suggest products you're likely to buy.	Average order value increases by 15–35%; customers are less likely to leave without buying.	<i>Amazon Personalize, Netflix</i>
AI-Generated Product Descriptions	Automatically writes unique, SEO-friendly product descriptions for thousands or millions of products.	Content creation time reduced by 90%; consistent brand voice across all products.	<i>Jasper, Shopify Magic</i>
Customer Behavior Prediction	Predicts whether a customer is likely to buy again, or is about to stop shopping with a brand.	Marketing teams can target the right customers at the right time, improving campaign ROI.	<i>Salesforce Einstein, Adobe Sensei</i>
Conversational Shopping Assistants	AI chatbots help customers find products by understanding natural language questions like 'show me something comfortable for a beach holiday'.	Conversion rates (how many visitors actually buy) increase by up to 25%.	<i>H&M Virtual Stylist, Sephora AI</i>
Demand Forecasting	Predicts how much of each product will be sold in a given period based on trends, seasons, and events.	Stockouts reduced by 30%; less money wasted on unsold inventory.	<i>Blue Yonder, Relex Solutions</i>
Dynamic Pricing Optimization	Adjusts prices automatically in real time based on demand, competitor prices, and available stock.	Maximizes revenue; stays competitive without manual price changes.	<i>Uber Surge Pricing model, Wiser</i>
AI-Powered Catalog Management	Automatically tags, categorizes, and describes products in a large product catalog.	Managing a catalog of millions of products becomes possible without a huge team.	<i>Algolia, Constructor.io</i>

Use Case	What GenAI Does	Business Impact	Real-World Example
Visual Search Systems	Lets customers upload a photo to find similar products — no need to describe what they want in words.	Customers can find products they couldn't describe; better discovery experience.	<i>Pinterest Lens, ASOS StyleMatch</i>
Review Summarization & Sentiment Analysis	Reads thousands of customer reviews and produces a short, balanced summary of what customers liked and didn't like.	Customers decide faster; brands identify product improvement areas more easily.	<i>Bazaarvoice, Yotpo</i>
Virtual Try-On Systems	Uses AI to digitally overlay clothing, glasses, or makeup on a customer's photo so they can 'try' before buying.	Product return rates reduced by up to 40% because customers are more confident in their purchase.	<i>Warby Parker, L'Oreal ModiFace</i>

5.2 Why Recommendations Are Such a Big Deal

I did some extra research on Amazon's recommendation engine because the numbers are so striking. About 35% of Amazon's total revenue is credited to its AI recommendation system. That is approximately \$70 billion in annual sales coming from a system that suggests products. For a student learning about AI, this is one of the clearest examples of how AI directly drives business value.

Modern GenAI has made recommendations even more powerful by adding conversational ability. Instead of just showing 'customers who bought this also bought that,' modern systems can understand a request like 'I need a gift for my mom who likes gardening and has a budget of 1000 rupees' and generate a personalized set of recommendations. That is a level of sophistication that was simply not possible a few years ago.

5.3 How These Use Cases Compare in ROI

Use Case	Reported Impact / Result
Personalized Recommendations	35% of Amazon's revenue; 15–35% increase in average order value
AI Product Descriptions	90% faster content production; better search rankings
Dynamic Pricing	2–5% revenue improvement; better margin management
Virtual Try-On	40% reduction in product returns
Demand Forecasting	30% fewer stockouts; 25% less excess inventory
Conversational Commerce	25% higher conversion rate; 40% fewer support calls

5.4 My Key Observations from Retail

Retail — What I Observed and Learned

- I was most surprised by the virtual try-on technology — the idea that AI can place a pair of glasses on your face from a photo and show how they look is genuinely impressive, and it directly solves the biggest problem in online fashion: not being able to touch or try products
- Amazon's recommendation engine contributing 35% of revenue was the most eye-opening statistic in this entire research — it shows how directly AI can be tied to business results
- AI-generated product descriptions are particularly useful for small businesses and startups that can't afford to hire large content writing teams
- Dynamic pricing needs to be used carefully — if customers notice prices change very frequently (like airline tickets), it can feel unfair and damage brand trust
- I personally experienced the Sephora AI makeup assistant and found it surprisingly accurate — this helped me understand how real these use cases have become

6. Information Technology & Software

As someone interested in data engineering and technology, this was the section I was most personally interested in. The IT and software industry is both creating GenAI tools and using them in its own work. The result is a kind of 'feedback loop' where AI is helping build better AI.

The most talked-about tool in this space is GitHub Copilot. A 2024 survey by GitHub found that 88% of developers who use Copilot say it makes them more productive, and developers complete tasks up to 55% faster with AI assistance. With over 1.8 million paid subscribers, it is the most widely adopted AI tool in software development today.

6.1 Use Case Overview Table

Use Case	What GenAI Does	Business Impact	Real-World Example
AI Code Generation	You describe what you want in plain English, and the AI writes the code for you. It also completes code as you type.	Development time reduced by 30–55%; developers spend more time on design and less on typing.	<i>GitHub Copilot, Amazon CodeWhisperer</i>
Automated Bug Detection & Resolution	AI scans your codebase, finds bugs and security issues, explains them, and suggests how to fix them.	QA cycles are shorter; expensive bugs caught earlier in development.	<i>DeepCode (Snyk), SonarQube AI</i>
DevOps Incident Management	When a system goes down or slows, AI reads all the logs and alerts to diagnose what went wrong and suggests a fix.	Time to fix production problems reduced from hours to minutes.	<i>PagerDuty AI, Dynatrace Davis</i>
AI-Powered Code Review	AI reviews new code submissions to check for logic errors, security risks, and style issues before a human reviews.	Code quality improves; senior developers spend less time on routine reviews.	<i>CodeRabbit, Sourcegraph Cody</i>
API Documentation Generation	Automatically writes API documentation, examples, and changelogs directly from the code.	No more out-of-date docs; developers save hours of documentation writing.	<i>Mintlify, Swimm AI</i>
Test Case Generation	AI reads your code and automatically writes tests to verify it works correctly.	Higher test coverage without the tedious work of writing tests manually.	<i>Diffblue Cover, Codium AI</i>
Infrastructure as Code (IaC) Automation	You describe what cloud infrastructure you need in plain English, and AI generates the configuration files.	Setting up cloud infrastructure takes minutes instead of days.	<i>Pulumi AI, AWS CloudFormation AI</i>
Security Threat Modeling	AI analyzes your system design and code to identify potential security	Security is built into the development process	<i>Microsoft Copilot for Security, Snyk</i>

Use Case	What GenAI Does	Business Impact	Real-World Example
	vulnerabilities and suggest how to fix them.	from the start, reducing expensive fixes later.	

6.2 What GitHub Copilot Is Teaching Us About AI Productivity

I found GitHub Copilot particularly interesting to research because there is a lot of real data about how it affects developer productivity. The key finding from GitHub's own research is that the biggest time saving is not actually in writing code faster — it is in reducing the time developers spend searching for answers, looking up documentation, and switching between tasks. This non-coding time is called 'context-switching,' and it can eat up 20–30% of a developer's day.

What this tells me is that AI tools are most valuable not when they replace a task entirely, but when they reduce the friction and interruptions in a human's workflow. That's an insight I want to remember as I continue learning about AI systems.

6.3 A Note on Security — AI Can Create New Risks Too

One thing I came across in my research that I think is important to mention: AI coding tools can also introduce new security risks. A 2022 study by Stanford University found that developers using AI coding assistants were actually more likely to have security vulnerabilities in their code — not because the AI wrote bad code, but because developers trusted the AI's suggestions without carefully reviewing them.

This is a good lesson: AI is a tool, not a replacement for human judgment. The best developers use AI to be faster and more productive, but they still review the output carefully. This idea of 'human in the loop' comes up across all industries, not just software.

Security Risk	How It Can Be Addressed
AI suggests code with security vulnerabilities	Always review AI-generated code; use AI security scanners like Snyk or Semgrep alongside coding AI
Developers trust AI output without checking	Build code review habits; never auto-accept AI suggestions in security-critical code
Hardcoded credentials in generated code	Configure AI tools to never suggest hardcoded passwords; use scanning tools to catch them
Prompt injection in AI-powered apps	Understand how LLMs can be manipulated through inputs; design applications with this in mind

6.4 My Key Observations from IT & Software

IT & Software — What I Observed and Learned
<ul style="list-style-type: none"> This section was most relevant to me personally as a data engineering intern — tools like GitHub Copilot and AI documentation generators are things I can start using right now

- The 55% productivity improvement from AI coding tools is real and measurable — this is not just marketing, it is backed by actual developer surveys and studies
- The Stanford security study was a good reminder that AI can also create new risks if used carelessly — I want to make sure I develop good habits around reviewing AI-generated code
- Infrastructure as Code automation particularly impressed me — the idea that you can describe cloud infrastructure in English and get working Terraform or Kubernetes configuration is something I find directly useful
- I think the most important shift in software is not any specific tool, but the change in how developers think — from 'I need to write this' to 'I need to design the right solution and guide AI to implement it'

7. Patterns I Noticed Across All Industries

After completing research on all five industries, I started to notice certain themes that kept repeating. These are not just coincidences — they seem to reflect what GenAI is fundamentally good at. Here are the five patterns I identified:

7.1 Pattern 1: Automating Documents and Reports

In every single industry I studied, one of the first and most impactful uses of GenAI was automating the creation of written documents. Clinical notes in healthcare, financial reports in banking, maintenance manuals in manufacturing, product descriptions in retail, API documentation in software — all of these involve a human expert converting raw data or observations into a written document.

GenAI is very good at this because that is essentially what language models are trained to do: produce structured, coherent text from input data. This is why document automation consistently appears as a high-ROI first use case across industries.

7.2 Pattern 2: Making Complex Systems Easier to Use

Across healthcare, banking, retail, and IT, I noticed that GenAI is being used to create natural language interfaces for complex systems. Instead of having to learn a special query language, fill out a complicated form, or navigate a multi-level menu, users can just describe what they want in normal language. This is making specialized systems accessible to people who would not normally be able to use them.

For example, a doctor who is not technically trained can ask an AI system 'show me all patients with high blood pressure who are also diabetic' instead of writing a database query. That is a meaningful improvement in usability.

7.3 Pattern 3: Predicting Problems Before They Happen

In all five industries, GenAI is being used to look at historical and real-time data and predict what is likely to happen next. Predicting machine failures in manufacturing, predicting loan defaults in banking, predicting product demand in retail, predicting patient health deterioration in healthcare, predicting software bugs in IT.

What struck me about this pattern is that it represents a fundamental shift in how organizations operate — from reacting to problems after they happen, to preventing them before they occur. That change has a compounding impact over time.

7.4 Pattern 4: Personalization at Large Scale

Every industry is using GenAI to offer personalized experiences to large numbers of people simultaneously. Healthcare gives personalized treatment recommendations. Banking gives personalized financial advice. Retail gives personalized product recommendations. Software gives personalized code suggestions.

What makes this interesting is that personalization used to require human expertise and time — a personal financial advisor, a personal shopper, a personal trainer. GenAI makes it possible to offer that kind of personalized attention to millions of people at once, at almost no extra cost.

7.5 Pattern 5: Helping Organizations Stay Compliant

Healthcare, banking, and manufacturing all operate under strict regulatory requirements. Staying compliant means reading through enormous volumes of rules, checking internal processes against them, and producing documentation that proves compliance. GenAI is being used to automate these compliance checks, which is reducing costs and reducing the risk of accidental violations.

7.6 Challenges I Found — Why GenAI Is Not Perfect

Along with these positive patterns, I also came across several recurring challenges that organizations face when deploying GenAI. These are important to understand because they explain why adoption is not happening overnight everywhere:

Challenge	What It Means in Simple Terms
Data Quality	AI is only as good as the data it is trained on. Bad data = bad AI outputs.
Hallucination	AI can confidently produce wrong answers. In high-stakes fields like medicine or banking, this is a serious risk.
Explainability	Many GenAI models cannot explain why they gave a particular output — this conflicts with regulations that require transparency.
Integration with Old Systems	Most organizations have old software systems that were not designed to work with AI. Connecting them is complex and expensive.
Talent Shortage	There are not enough people who understand how to build, deploy, and manage GenAI systems properly.

8. Future Trends — What I Think Is Coming Next

In the final part of my research, I looked at what experts and companies are predicting for the future of GenAI. Here are five trends that I found most interesting and believable:

8.1 Agentic AI — AI That Takes Actions on Its Own

Right now, most GenAI tools respond to a question or request and give you an answer. The next step — called Agentic AI — is systems that can break a complex goal into steps, take actions (like searching the web, running code, or filling out forms), and complete a task with minimal human intervention.

For example, instead of asking an AI 'how do I fix this bug,' an agentic AI system would read the bug report, search the codebase, write a fix, test it, and open a pull request — all by itself. Companies like Cognition (which built a product called Devin) are already demonstrating early versions of this.

In manufacturing, an agentic AI could monitor production, detect a problem, identify the root cause, adjust machine settings, and update the maintenance log — all automatically, overnight.

8.2 Multimodal AI — AI That Understands Everything at Once

Most current AI systems are good at one type of data — either text, or images, or audio. Multimodal AI systems can process all of these together at the same time. For healthcare, this means an AI that simultaneously reads a patient's text medical history, looks at their scan images, listens to their voice symptoms, and produces a unified recommendation. We are already seeing early examples of this with GPT-4 Vision and Google Gemini.

8.3 Domain-Specific AI Models

Rather than using one general AI model for everything, organizations are beginning to train specialized AI models on industry-specific data. Examples I came across:

Industry	Specialized AI Model Being Developed
Medicine	Med-PaLM 2 (Google), BioMedLM (Stanford) — trained on medical literature and clinical notes
Finance	BloombergGPT — trained on 700 billion tokens of financial data, news, and market information
Law	Harvey AI, CoCounsel — trained on legal case data and regulatory documents
Science	Galactica (Meta) — trained on research papers and scientific knowledge
Software Engineering	Code Llama, StarCoder, DeepSeek Coder — trained primarily on code repositories

8.4 Better Collaboration Between Humans and AI

I think one of the most important future developments is not a specific technology, but a change in how people and organizations think about working with AI. The most effective use of GenAI is not to replace humans entirely, but to split tasks smartly — AI handles the high-volume, repetitive, or data-intensive parts, while humans focus on judgment, creativity, relationships, and ethical decisions.

Some of the best ways to think about this split:

- AI as a first drafter — AI writes a first version; a human reviews, corrects, and approves it
- AI as an analyst — AI finds patterns in data; a human decides what to do with those insights
- AI as a quality checker — AI reviews human work for errors before it goes out
- AI as a coach — AI gives real-time suggestions and guidance while a human does their job

8.5 Real-Time AI Everywhere

As AI models get smaller and cheaper to run, they will start to appear in more and more everyday tools and devices. Already, AI is available directly inside code editors, document software, and email tools. In the future, it will be embedded in medical devices in operating theatres, factory floor sensors and displays, trading terminals in financial markets, and even in everyday consumer devices. The goal is for AI assistance to be available exactly when and where you need it.

9. Conclusion and Personal Learnings

This Week 01 research project has been a genuinely eye-opening experience for me. Before I started, I knew that AI was being used in many industries, but I did not appreciate the breadth and depth of real, deployed applications that are already changing how work gets done across healthcare, finance, manufacturing, retail, and technology.

The 47 use cases I documented in this report are not hypothetical or experimental — they are real products being used by real companies right now. That fact changed my perspective significantly.

9.1 My Top 5 Personal Learnings

5. Documentation automation is the most universally applicable use case — every industry has knowledge workers who spend large amounts of time writing structured documents, and GenAI can help all of them.
6. Generative AI is not just about language. The same core technology is being used to discover drug molecules, simulate factory floors, analyze satellite imagery for retail demand forecasting, and generate software code. Its applications are much broader than just chatbots.
7. Hallucination is the biggest technical challenge — the fact that AI can confidently produce wrong answers is a fundamental limitation that affects how much trust we can place in AI outputs, especially in high-stakes settings.
8. Human oversight is not optional — in every high-stakes use case I studied (medical AI, financial AI, legal AI), the best implementations keep a human in the review loop. AI speeds up the work; humans provide the judgment and accountability.
9. As a data engineering student, the use cases in IT & Software are most directly relevant to me. Tools like GitHub Copilot, AI documentation generators, and IaC automation are things I can start learning and using today.

9.2 What I Would Like to Research Further

Based on this week's research, here are the areas I am most interested in exploring in the coming weeks:

- How Retrieval-Augmented Generation (RAG) works and how it reduces hallucination in enterprise AI systems
- How to evaluate whether a GenAI model is actually performing well on a specific task
- How organizations set up AI governance — policies that ensure AI is used responsibly
- Practical implementation of GenAI in data engineering workflows, which is most relevant to my internship focus

Summary of This Report

- 5 industries analyzed: Healthcare, Banking, Manufacturing, Retail, IT & Software
- 47 real-world use cases documented with explanations, impacts, and company examples
- 5 cross-industry patterns identified
- 5 key implementation challenges explained

- 5 future trends analyzed
- 25 references from academic papers, industry reports, and official sources
- Submitted as Week 01 deliverable for MacroEdtech GenAI Research Internship — June 2026

End of Week 01 Research Document — Shivya | MacroEdtech GenAI Research Internship | June 2026

Appendix A: GenAI Tools I Explored During Research

While researching for this document, I explored several GenAI tools personally. Here is a summary of the main tools I came across and their primary use:

Tool / Platform	What It Does and Why It Stood Out to Me
ChatGPT (OpenAI)	General-purpose LLM; I used it during research to clarify concepts I didn't understand. Impressive at summarizing complex topics.
GitHub Copilot	AI coding assistant integrated into VS Code. Completes code as you type. Most practical tool for developers.
Nuance DAX (Microsoft)	Healthcare documentation AI. Generates clinical notes from doctor-patient conversations. Very specialized.
Elicit	Research assistant that summarizes academic papers. Extremely useful for research tasks like this one.
Midjourney / DALL-E	Image generation tools. I explored these to understand diffusion models — less directly relevant to enterprise use cases.
BloombergGPT	Finance-specific LLM trained on Bloomberg data. Shows how domain-specific models outperform general models in specialized fields.
NVIDIA Omniverse	Digital twin platform for industrial use. The factory simulation demos available online are remarkable.
Shopify Magic	AI tool that generates product descriptions for e-commerce stores. Very practical retail AI application.

Appendix B: Key Terms and Definitions (Glossary)

Here are the most important technical terms I came across during this research, explained in simple language:

Term	Simple Definition
Generative AI (GenAI)	AI that creates new content (text, images, code, audio) rather than just classifying or analyzing existing data.
Large Language Model (LLM)	A very large AI model trained on huge amounts of text that can read, write, and understand language at a high level.
Hallucination	When AI confidently generates incorrect or made-up information. A key risk in real-world deployments.
Prompt Engineering	Writing clear and specific instructions to get the best results from an AI system.
Fine-tuning	Training an existing AI model further on specific data to make it better at a particular task.
RAG (Retrieval-Augmented Generation)	A method where the AI retrieves relevant real documents before generating an answer, reducing hallucination.
Transformer Architecture	The mathematical design used in most modern AI language models. Introduced in 2017 and still the foundation of GPT, Gemini, Claude, etc.
Digital Twin	A virtual, real-time copy of a physical machine or factory used for testing and monitoring.
Agentic AI	AI systems that can plan, take actions, and complete multi-step tasks without much human guidance.
RLHF	Reinforcement Learning from Human Feedback — how AI models are taught to give better answers using ratings from humans.
Foundation Model	A very large AI model trained on broad data that can be used for many different tasks.
Multimodal AI	AI that can process and understand text, images, audio, and other types of data together.
Inference	Running a trained AI model to get an answer from new input — the 'using' stage of AI, as opposed to the 'training' stage.

Term	Simple Definition
Zero-shot Learning	The ability of an AI model to do a task it was never specifically trained on, just from a plain-language description.
IaC (Infrastructure as Code)	Writing code that describes cloud infrastructure (servers, networks, databases) so it can be automatically set up.

10. References

The following references were consulted while preparing this research document. References are listed in APA 7th Edition format.

10.1 Academic Papers

- [1] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30. <https://arxiv.org/abs/1706.03762>
- [2] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901. <https://arxiv.org/abs/2005.14165>
- [3] Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>
- [4] Chen, M., Tworek, J., Jun, H., Yuan, Q., Kaplan, J., & Zaremba, W. (2021). Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*. <https://arxiv.org/abs/2107.03374>
- [5] Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., & Ting, D. S. W. (2023). Large language models in medicine. *Nature Medicine*, 29(8), 1930–1940. <https://doi.org/10.1038/s41591-023-02448-8>
- [6] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., & Lowe, R. (2022). Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*. <https://arxiv.org/abs/2203.02155>
- [7] Pearce, H., Ahmad, B., Tan, B., Dolan-Gavitt, B., & Karri, R. (2022). Asleep at the keyboard? Assessing the security of GitHub Copilot's code contributions. *IEEE Symposium on Security and Privacy (SP)*, 754–768. <https://doi.org/10.1109/SP46214.2022.9833571>
- [8] Bommasani, R., Hudson, D. A., Altman, R., Arora, S., & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*. <https://arxiv.org/abs/2108.07258>

10.2 Industry Reports

- [9] McKinsey Global Institute. (2023). The economic potential of generative AI: The next productivity frontier. McKinsey & Company. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai>
- [10] McKinsey & Company. (2024). Generative AI in healthcare: Opportunities and challenges. McKinsey Health Institute. <https://www.mckinsey.com/industries/healthcare>
- [11] Goldman Sachs. (2023). Generative AI: Too much spend, too little benefit? Goldman Sachs Research. <https://www.goldmansachs.com/intelligence/pages/generative-ai-too-much-spend.html>
- [12] World Economic Forum. (2023). Future of Jobs Report 2023. <https://www.weforum.org/reports/the-future-of-jobs-report-2023>
- [13] Gartner. (2024). Top strategic technology trends for 2024. <https://www.gartner.com/en/articles/gartner-top-10-strategic-technology-trends-for-2024>
- [14] Salesforce. (2024). State of AI in retail and consumer goods. <https://www.salesforce.com/resources/research-reports/state-of-ai/>
- [15] Association of Certified Fraud Examiners. (2024). Report to the nations: 2024 global study on occupational fraud and abuse. <https://www.acfe.com/report-to-the-nations/2024/>

- [16] GitHub. (2024). GitHub Copilot enterprise features and research insights. GitHub Blog. <https://github.blog/2024-02-27-github-copilot-enterprise-is-now-generally-available/>

10.3 Company and Product Documentation

- [17] Microsoft. (2024). Nuance DAX Copilot clinical documentation AI. <https://www.nuance.com/healthcare/ambient-clinical-intelligence.html>
- [18] OpenAI. (2023). GPT-4 technical report. arXiv:2303.08774. <https://arxiv.org/abs/2303.08774>
- [19] Anthropic. (2024). Claude model card and documentation. <https://www.anthropic.com/claude>
- [20] NVIDIA. (2024). NVIDIA Omniverse for industrial digital twins. <https://developer.nvidia.com/omniverse>
- [21] Siemens. (2024). MindSphere digital manufacturing platform overview. <https://www.siemens.com/global/en/products/automation/topic-areas/digital-enterprise/>
- [22] Insilico Medicine. (2023). First fully AI-designed drug enters Phase II clinical trial. <https://insilico.com>

10.4 Regulatory and Governance Sources

- [23] U.S. Food and Drug Administration. (2023). Artificial intelligence and machine learning enabled medical devices. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices>
- [24] European Parliament. (2024). Regulation on artificial intelligence (EU AI Act). <https://artificialintelligenceact.eu/>
- [25] Federal Reserve System. (2011, updated 2021). Supervisory guidance on model risk management (SR 11-7). <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>

End of References