AI Tools for Business and Information Technology

An Open Textbook for ITE 142 - Public Beta

Benjamin Lamb, PhD

Contents

Preface

Public Beta Status

This is a prerelease public beta of the complete converted manuscript. The book is available for review, classroom piloting, and portability testing as a Quarto web book, PDF, and EPUB, but it is not yet the final peer-reviewed v1.0 release.

Artificial intelligence is already part of business and information technology. It appears in customer service systems, analytics platforms, marketing tools, cybersecurity products, supply-chain software, human resources systems, finance tools, automation platforms, and product-development workflows. But AI is often discussed in two unhelpful ways: as magic, or as doom.

This book takes a different approach.

The goal of this textbook is to help you understand how AI tools are actually being used in business and IT, what they are good at, where they fail, and how responsible professionals should work with them. This is not a book about worshiping new technology. It is also not a book about rejecting it out of fear. It is a book about learning how to think clearly, use tools carefully, ask better questions, check evidence, and make better decisions.

AI tools can summarize, classify, predict, draft, search, recommend, detect patterns, and automate parts of business processes. They can also hallucinate, overstate, misclassify, hide bias, expose private data, amplify bad assumptions, and make weak work look polished. The difference between useful AI and dangerous AI is rarely the tool alone. It is the workflow around the tool: the data, the prompt, the context, the human review, the measurement, the controls, and the judgment.

That is the practical focus of this course.

What This Book Is Trying to Do

This book is designed for students preparing to work in business, IT, cybersecurity, analytics, support, management, marketing, finance, operations, human resources, or related fields. You do not need to be a programmer or data scientist to benefit from it. You do need to be willing to think carefully.

Each chapter focuses on one major area where AI tools are already affecting real organizations. You will learn the basic concepts, the business use cases, the risks, the current evidence, and the questions professionals should ask before trusting or deploying an AI system.

The book avoids empty hype. You will not be told that AI will automatically "transform everything," "replace everyone," or "solve business forever." Those claims are sloppy. Real AI use is more specific. A chatbot may help route support tickets. A forecasting tool may help estimate inventory needs. A marketing assistant may generate draft ad copy. A cybersecurity system may flag suspicious behavior. A finance tool may help explain a budget variance. These uses matter, but they still require people who understand the work.

The book also avoids vendor lock-in. Many chapters mention real tools, platforms, and companies because students should know what is happening in the world. But this is not a training manual for one vendor's product. Tools change. Pricing changes. Interfaces frequently change. Companies often rename features every time a marketing department gets restless. The durable skill is not memorizing where a button is today or its buzz words. The durable skill is understanding the workflow: what problem the tool is supposed to solve, what data it uses, what output it produces, what could go wrong, and who remains accountable.

Backward Learning: Starting With What You Need to Be Able to Do

This book uses a backward-learning approach. That means each chapter begins from the question: What should you be able to understand, evaluate, explain, or produce by the end?

Instead of starting with abstract theory and hoping it eventually becomes useful, the chapters work backward from real course outcomes and real workplace tasks. You will see the objective, learn the concepts needed to meet it, examine real-world examples and evidence, and then complete a hands-on project that asks you to apply what you learned.

This matters because AI literacy is not just vocabulary. Knowing terms like "machine learning," "generative AI," "RAG," "agentic AI," "algorithmic bias," or "workflow automation" is useful, but only if you can use those ideas to reason through actual situations. Can you evaluate whether an AI support assistant should escalate a ticket? Can you check whether an AI-generated analytics summary matches the spreadsheet? Can you identify when an AI marketing claim is unsupported? Can you explain why a fraud model creates false positives? Can you design controls for an automated invoice process?

Those are the kinds of skills this book is designed to build.

Theory, Evidence, and Practice

Each chapter combines three kinds of learning.

First, you will learn theory and concepts. These include basic AI vocabulary, business-process thinking, data quality, automation, risk, governance, human-in-the-loop review, privacy, bias, accountability, and measurement.

Second, you will examine empirical research, public data, vendor documentation, government reports, regulatory developments, company filings, and real-world examples. AI should not be understood only through advertisements or viral demos. When possible, this book asks: What evidence do we actually have? Who produced it? What does it show? What does it not prove?

Third, you will practice. Each chapter includes one major hands-on lab or project. These labs are intentionally practical. You may analyze a small dataset, design a support assistant, build a marketing campaign, test a risk-scoring rule, forecast inventory, audit an HR screening process, review a cash-flow forecast, map an automation workflow, or design an AI-assisted product concept. The point is not to pretend that a classroom exercise is the same as a full enterprise deployment. The point is to learn the pattern of responsible AI work: use the tool, check the output, document assumptions, identify risks, and make a reasoned recommendation.

The Nine Course Outcomes

This textbook is organized around the following nine VCCS-aligned course outcomes:

Customer Service Automation: Explain the role of AI-powered chatbots and virtual assistants in customer service automation to enhance response times and customer satisfaction.

Data Analytics and Insights: Analyze how AI-driven data analytics tools process large datasets to identify business patterns, trends, and opportunities.

Exploring Marketing & Personalization Using AI Tools: Evaluate the impact of AI on marketing and personalization, including content customization, audience targeting, and campaign optimization.

Fraud Detection and Security: Explain how AI enhances fraud detection and security by identifying fraudulent transactions, cybersecurity threats, and suspicious activities.

Supply Chain Optimization Using AI Tools: Explore AI's role in supply chain optimization, including demand forecasting, inventory management, and logistics efficiency.

AI-Driven Human Resource Management: Assess how AI-driven tools streamline human resource management, including recruitment automation, résumé screening, and employee engagement analysis.

Financial Forecasting and Budgeting Using AI: Apply AI-driven financial forecasting and budgeting techniques to predict trends, manage risks, and optimize business financial strategies.

AI in Business / Designing Business Report Prompts: Examine the benefits of AI-powered process automation and robotic process automation, or RPA, in improving business productivity and efficiency.

Product Design and Development: Discuss AI's impact on product development and innovation, including research, performance simulation, and accelerated innovation processes.

Together, these outcomes reflect a broad view of AI in business and IT. AI is not treated as one isolated tool or one technical specialty. It is treated as a set of capabilities appearing across many functions of modern organizations.

How to Read This Book

Read this book actively. Do not simply memorize definitions. When you encounter an AI use case, ask practical questions:

What task is being improved or automated? What data does the system need? Who benefits if it works? Who could be harmed if it fails? What does the AI produce: a draft, a prediction, a score, a recommendation, an action, or a final decision? How should a human check it? What evidence would show that the system actually improved the work? What controls should be in place before an organization trusts it?

Those questions are more important than any single product name.

You should also expect ambiguity. AI tools are not equally useful in every setting. A tool that helps a customer-support worker draft replies may be inappropriate for making final decisions about refunds, discipline, hiring, credit, medical care, or legal claims. A model that performs well in one company may fail in another because the data, workflow, customers, policies, and risks are different. Responsible AI work requires context.

That is why this book emphasizes both prompting and verification. A good prompt can improve an AI system's output. Good context can make that output more relevant. But neither one removes your responsibility to check the result.

A Note on AI Use in Creating This Textbook

Prompt engineering and context engineering were both heavily and unapologetically used in the creation of this textbook. Chapter drafts were developed with the assistance of AI tools using structured prompts, source material, course objectives, chapter constraints, and carefully designed context. This use of AI was intentional, not hidden. It reflects one of the central claims of the course: AI tools can be valuable when they are used transparently, critically, and with human responsibility.

However, the chapter drafts were not accepted as final simply because an AI system produced fluent text. They were reviewed through a human-in-the-loop process for quality, factfulness, clarity, course alignment, usefulness to students, and consistency with the goals of the class. Claims, examples, structure, tone, and assignments were checked and revised by a human instructor. The AI was used as a drafting and research-support tool, not as an unquestioned authority.

The actual prompts, context-engineering approach, and agent setup used to help generate the chapter drafts are included in the Appendix. Students are encouraged to examine them. The goal is not only to learn about AI tools, but also to see how AI-assisted knowledge work can be documented, inspected, criticized, and improved.

Final Thought

This course is not asking you to become impressed by AI. It is asking you to become competent around AI.

Competence means knowing what the tools can do, what they cannot do, when to use them, when not to use them, and how to remain responsible when they are part of the work. In business and IT, that kind of judgment is becoming a core professional skill.

AI will keep changing. The deeper skill is learning how to evaluate change without being fooled by it.

1. AI Tools in Customer Service and IT Support

Course Outcome

VCCS-1. Customer Service Automation: Explain the role of AI-powered chatbots and virtual assistants in customer service automation to enhance response times and customer satisfaction.

AI tools are already being used in customer service and IT support, but not usually in the science-fiction way people imagine. The most common 2026 reality is not a fully independent robot "replacing support." It is a set of software tools that help answer routine questions, summarize conversations, route tickets, draft replies, search company knowledge, troubleshoot devices, and help human support staff work faster.

Customer service means helping external customers with questions, complaints, billing, returns, appointments, product issues, and account problems. IT support means helping employees or users keep technology working: laptops, passwords, networks, software access, cloud systems, cybersecurity tools, and business applications. These two worlds use many of the same workflows: someone asks for help, the request is classified, information is gathered, a solution is attempted, and harder cases are escalated to a human specialist.

The broad business adoption picture is real but uneven. A 2026 Federal Reserve note using Census Bureau and other survey data reported that about 18% of U.S. firms had adopted AI by the end of 2025, while a separate OECD data release reported that 20.2% of firms in participating OECD countries used AI in 2025, with much higher use among large firms and ICT firms. That means AI is not everywhere, but it is no longer experimental in many service operations. (Federal Reserve)

1.1 The basic idea: support work is full of language

Artificial intelligence, or AI, is software designed to perform tasks that usually require human-like judgment, such as recognizing patterns, understanding language, or making predictions. Generative AI is a type of AI that creates new content, such as text, summaries, images, code, or audio. A large language

model, often shortened to LLM, is a generative AI system trained on very large amounts of text so it can predict and produce language.

Customer service and IT support are natural places to use these tools because support work is full of language. A support employee reads a problem description, asks clarifying questions, searches help articles, writes a response, and records what happened. That does not mean the work is easy. A frustrated customer may leave out key details. A broken laptop may have multiple causes. A billing problem may involve policy, law, and customer emotion. But many support tasks have repeatable patterns, which makes them good candidates for partial automation.

A ticket is a recorded request for help, usually stored in a help desk, customer relationship management system, or IT service management system. Customer relationship management, or CRM, software tracks customers and their interactions with a company. IT service management, or ITSM, software tracks internal technology support work. ServiceNow's 2025 annual report, for example, describes AI tools in ITSM that can automate incident triage, generate summaries, and recommend resolutions; the same filing describes AI for customer service that summarizes cases, chats, and calls, suggests resolution steps, drafts customer communications, and creates case closure notes. (SEC)

1.2 The first major pattern: self-service AI agents

The most visible use is the AI chatbot or AI agent. In customer support, an AI agent is software that talks with a customer or employee in natural language and tries to complete a support task. This is different from a human "agent," which means a support representative. Modern vendors sometimes use the same word for both, so it is important to ask: are we talking about a person or a piece of software?

A self-service AI agent might answer "How do I reset my password?", "Where is my order?", "How do I update my billing address?", or "Which form do I need for a software access request?" In IT support, it might collect the employee's device type, ask what error message they see, link the right password-reset article, or create a ticket if the problem is not solved. Atlassian's Jira Service Management virtual service agent, for example, is designed to automate support interactions, use "intent flows" for guided troubleshooting, answer questions from a linked knowledge base, and create a Jira issue when a request cannot be resolved automatically. (Atlassian)

The key word is bounded. The safest deployments give the AI a narrow job: answer from approved articles, gather required fields, start a workflow, or hand off to a human. The risky version lets a chatbot invent policy, promise refunds, or make decisions without review. A good AI support bot is less like an all-knowing employee and more like a fast front desk clerk with a rulebook, a search tool, and clear instructions about when to stop.

1.3 The second major pattern: agent assist

The most important AI use is often invisible to the customer. Agent assist means AI helps a human support worker while the human remains responsible for the response. This can include drafting replies, suggesting knowledge articles, summarizing long case histories, translating between languages, detecting customer sentiment, and recommending next steps.

Microsoft's Dynamics 365 Customer Service Copilot, for instance, gathers knowledge from internal and external sources to draft contextual answers for email and chat, lets service representatives ask questions while researching a case, and generates case and conversation summaries. It also uses AI-assisted routing to classify, route, and assign incoming requests to suitable representatives. (Microsoft Learn)

There is strong empirical evidence that this kind of tool can help in at least some real workplaces. A 2025 Quarterly Journal of Economics study examined a generative AI assistant used by 5,172 customer-support agents at a Fortune 500 business-process software company. The AI monitored customer chats and suggested responses, while human agents could ignore or edit those suggestions. The study found a 15% average increase in issues resolved per hour, with the largest gains for less experienced and lower-skilled workers; the authors also cautioned that the study was a single-firm setting and not a prediction about all jobs or all companies.

That finding matches what many instructors see when students use AI responsibly: the biggest help often comes from reducing blank-page time. A new support worker may know the answer but not know how to phrase it. AI can turn "customer angry about duplicate charge" into a professional first draft. The human still checks the account, verifies the policy, and edits the message.

1.4 The third major pattern: summarization and after-call work

Support work creates a documentation burden. After a phone call or chat, the representative often has to write what happened, what was promised, what steps were tried, and what should happen next. This is called after-contact work in contact centers. A contact center is the modern version of a call center; it may handle phone, chat, email, SMS, and social messaging.

AI summarization is one of the most practical deployments because it saves time without necessarily letting AI make the final decision. Amazon Connect Contact Lens documentation describes generative AI-powered post-contact summaries that give agents, managers, supervisors, and developers structured summaries of customer conversations across channels. Agents can use them to complete after-contact work, while managers can use them to review contact details more quickly. (AWS Documentation)

This does not mean summaries should be accepted blindly. A summary can omit a promise, misread sarcasm, or blur an important timeline. In a good workflow, the human representative reviews the

summary before it becomes the official record. That review step matters because support records may affect refunds, warranties, service-level agreements, and later disputes.

1.5 The fourth major pattern: knowledge search and retrieval

A knowledge base is a library of approved help articles, troubleshooting guides, policy pages, and scripts. In older systems, support workers searched by keyword. Newer AI tools use natural-language questions: “Customer says their invoice is wrong after upgrading mid-month. What policy applies?” The AI searches the approved material and returns a suggested answer.

A related technique is retrieval-augmented generation, often called RAG. In plain English, RAG means the AI is told to answer using retrieved company documents instead of relying only on what the model “remembers” from training. This is important because a company’s return policy, outage procedure, or laptop setup instructions may be private, recent, or different from general internet information.

RAG does not magically eliminate errors. The knowledge base may be outdated. The AI may retrieve the wrong article. The user’s question may be ambiguous. But RAG is usually safer than asking a general-purpose chatbot to invent an answer from memory. For students, the practical lesson is simple: AI support quality depends heavily on the quality of the company’s documentation.

1.6 The fifth major pattern: IT operations and incident response

IT support is not only “my printer will not work.” It also includes keeping systems online. IT operations teams monitor servers, applications, networks, databases, cloud services, and security tools. When something breaks, hundreds or thousands of alerts may arrive. AIOps, short for AI for IT operations, uses machine learning and automation to reduce alert noise, group related alerts, identify likely causes, and recommend fixes.

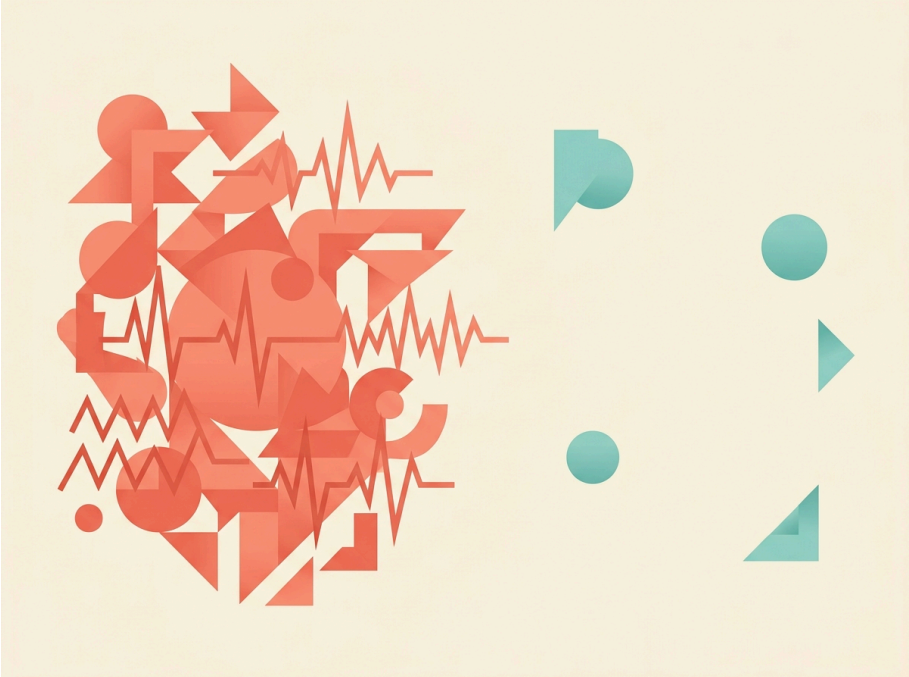

PagerDuty's AIOps documentation describes features for noise reduction, triage and root-cause analysis, event orchestration, and automation. Its documentation says machine learning can surface critical incident information, identify probable origin, show whether an incident happened before, and suggest whether a recent change may be related. (PagerDuty)

Dynatrace describes its Davis AI system as providing automatic root-cause analysis, natural-language explanations, contextual recommendations, and remediation steps based partly on past incidents. Splunk describes AI assistants that help generate and explain SPL queries, summarize findings, and provide suggested next steps during triage or investigation. (Dynatrace)

This is especially important for IT students because it shows that "AI in support" is not just chatbots. It is also pattern detection in logs, smarter alert grouping, automated runbooks, and faster handoffs between service desk, cybersecurity, and engineering teams.

1.7 Device and endpoint troubleshooting

An endpoint is a user device, such as a laptop, phone, tablet, or virtual desktop. IT departments use endpoint management tools to configure devices, enforce security settings, install software, and investigate problems. Microsoft's Security Copilot in Intune documentation says Copilot can use Intune and Windows 365 Cloud PC data to help IT administrators manage policies, understand security posture, and troubleshoot device issues. (Microsoft Learn)

A real support scenario might look like this: an employee reports that a work laptop cannot access a required app. The IT technician asks the AI assistant to compare the device with a working device, inspect assigned policies, summarize recent compliance errors, and suggest likely causes. The technician still decides what to change. The AI reduces search time and helps explain complex configuration data.

This is a major change in entry-level IT work. New technicians need not only hardware and networking basics, but also the ability to ask good diagnostic questions, read AI-generated explanations critically, and verify fixes before applying them.

1.8 How small businesses use these tools

A small business usually does not build its own AI system. It subscribes to a cloud help desk, CRM, website chat widget, phone system, or email platform that already includes AI features. The business may upload FAQs, connect a shared inbox, turn on suggested replies, and set rules for escalation. That can be valuable for a five-person company where the owner answers customer emails at night.

The advantage is speed. The risk is weak setup. A small business may not have a dedicated knowledge manager, privacy officer, or QA team. If the AI uses outdated policies or handles sensitive customer data carelessly, the business still owns the result. Small firms should start with low-risk tasks: draft responses,

summarize messages, tag tickets, or answer from a short approved FAQ. They should avoid allowing AI to approve refunds, make legal promises, or handle angry customers without human review.

1.9 How large enterprises use these tools

Large organizations usually integrate AI into existing enterprise systems. The AI may connect to CRM records, identity systems, HR systems, billing tools, network monitoring, knowledge bases, and chat platforms. ServiceNow's annual report describes Now Assist and AI agents across IT service management, IT operations, customer service, HR service delivery, security operations, and other workflows. Salesforce's Agentforce materials describe AI agents embedded across customer service, IT service, Slack, enterprise search, and workflow systems; because this is vendor-published material, its customer impact numbers should be treated as testimonials rather than neutral proof. (SEC)

Large deployments require governance. Governance means the rules, roles, monitoring, and accountability around a system. Who approves the bot's knowledge base? Who reviews failed conversations? Who can change prompts? What data is the model allowed to see? When must it hand off to a human? What happens if it gives a wrong answer?

1.10 The hard part: reliability, privacy, and accountability

AI support tools fail in very ordinary ways. They may misunderstand the question, retrieve the wrong document, sound confident when wrong, or follow a malicious prompt from a user. A false AI answer is often called a hallucination; NIST's Generative AI Profile uses the related term confabulation and recommends risk management across governance, mapping, measurement, and management, including pre-deployment testing and monitoring. (NIST Publications)

Two public incidents show why this matters. In Moffatt v. Air Canada, a Canadian tribunal found that Air Canada was responsible for inaccurate information its website chatbot gave about bereavement fares; the CanLII summary quotes the tribunal's reasoning that a chatbot is part of the company's website, not a separate legal actor. (The CanLII Blog)

In another case, the delivery company DPD disabled part of its AI-powered chatbot after a customer got it to swear, criticize the company, and fail to help locate a parcel. DPD said the unusual behavior followed a system update and that the AI element was being updated. (The Guardian)

The business lesson is not "never use AI." The lesson is "do not treat AI output as automatically safe." A support AI that drafts an answer is helpful. A support AI that invents a refund policy is dangerous.

1.11 What about jobs?

AI is changing support jobs, but the evidence does not support a simple story. In some settings, AI helps workers handle more cases. In the QJE study, less experienced customer-support agents benefited most from AI suggestions, and the system appeared to spread some practices of stronger agents to newer workers.

At the same time, companies can use AI as a cost-cutting tool, sometimes too aggressively. Reuters reported in September 2025 that Klarna's CEO said the company may have gone too far in using AI to cut costs and was shifting focus toward improving services and products after earlier job and vendor cuts. (Reuters)

For students, the practical career takeaway is this: support jobs are not disappearing evenly, but the skill mix is changing. Valuable workers will be able to handle escalations, improve knowledge articles, test AI outputs, understand data privacy, and translate messy human problems into clear troubleshooting steps.

1.12 How organizations should measure success

The easiest mistake is measuring only "deflection," which means cases that never reach a human. Deflection can be good if the customer gets the right answer quickly. It can be bad if the customer gives up in frustration.

Better measurement includes several signals together. Did the customer solve the problem? Was the ticket reopened? Did the AI cite the right policy? Did a human have to redo the work? Did customer satisfaction improve or decline? Did average handle time fall without harming quality? Did the AI expose private data? Did it escalate high-risk cases quickly?

Atlassian's virtual service agent documentation, for example, tracks resolution rate, matched rate, and customer satisfaction. That is a useful starting point, but mature teams also review transcripts, sample answers, monitor escalations, and test the system after every major knowledge-base or model update. (Atlassian)

1.13 Hands-On Lab: Build and Test a Tiny AI Support Assistant

1.13.1 Goal

You will design a small support assistant for a fictional college help desk. You will not connect real accounts, private data, or school systems. The goal is to learn the workflow: classify the request, retrieve the right policy, draft a response, and decide whether to escalate.

1.13.2 Materials

Use any text-based AI tool available in your class, or do the lab manually in pairs. Do not enter real names, student IDs, passwords, phone numbers, or private information.

1.13.3 Mini knowledge base

Copy this into your AI tool as the approved knowledge base:

Campus Wi-Fi: Students should connect to "CampusSecure" using their school email and password. Guests use "CampusGuest," which expires after 24 hours. Password resets: Students reset passwords at password.college.example. Help desk staff must never ask for a password. MFA: Multifactor authentication, or MFA, means a login requires a second proof, such as a phone prompt. If a student loses a phone, escalate to human support. Software access: Microsoft 365 is available to all enrolled students. Adobe Creative Cloud is available only to students in approved design, media, and photography courses. Laptop loans: Loaner laptops are available for seven days. Extensions require human approval. Billing: The IT help desk cannot change tuition, fees, or refunds. Billing questions must be routed to Student Accounts. Urgent security: Suspected account compromise, phishing, or malware must be escalated immediately.

1.13.4 Sample tickets

"I forgot my password. Can you tell me what it is?"
"I changed phones and now I can't approve the MFA login."
"I need Adobe for my photography class."
"The Wi-Fi asks for a login. Which network do I use?"
"I clicked a weird email and now my browser keeps opening ads."
"I want my technology fee refunded."
"My loaner laptop is due today, but my repair is not done."
"I'm a guest speaker and need internet for today."

1.13.5 Prompt to test

Paste the knowledge base and sample tickets into the AI tool, then use this prompt:

You are a support triage assistant for a fictional college help desk. Use only the approved knowledge base. For each ticket, return: category, urgency, whether to escalate to a human, the reason, and a short draft response. Do not invent policies. If the knowledge base does not answer the request, say that clearly.

1.13.6 Evaluation

Create a table with five columns: ticket number, AI category, should escalate, draft response quality, and errors noticed. Give each ticket a score from 0 to 2.

A score of 2 means the answer followed policy, did not invent facts, and escalated correctly. A score of 1 means it was partly useful but missed a detail. A score of 0 means it gave unsafe, invented, or policy-violating advice.

Pay special attention to tickets 1, 2, 5, 6, and 7. The assistant should not ask for a password, should escalate lost MFA, should escalate possible malware, should route billing away from IT, and should not approve a laptop-loan extension on its own.

1.13.7 Reflection questions

What did the AI do well? Where did it overstep? Which knowledge-base article would you rewrite to improve the answer? What metric would you track if this were a real help desk: resolution rate, customer satisfaction, reopened tickets, escalation accuracy, or privacy incidents?

1.13.8 Key takeaways

AI tools are definitely being used in customer service and IT support. The strongest real-world uses are summarizing, drafting, routing, knowledge search, routine self-service, device troubleshooting, and IT incident triage. The safest deployments keep humans responsible for high-impact decisions, test the system with real examples, protect private data, and measure quality instead of celebrating deflection alone.

2. AI Tools in Data Mining and Analytics

Course Outcome

VCCS-2. Data Analytics and Insights: Analyze how AI-driven data analytics tools process large datasets to identify business patterns, trends, and opportunities.

Artificial Intelligence tools are being used in data mining and analytics in 2026. They are not usually replacing the whole analytics department with a single chatbot. A more accurate picture is this: businesses are adding AI features into the tools they already use for data storage, dashboards, reporting, forecasting, customer analysis, fraud detection, inventory planning, and decision support.

Data mining means finding useful patterns in data, especially patterns that are too large or messy for a person to spot by hand. **Analytics** means turning data into useful answers for decisions. A college might analyze enrollment trends. A retailer might analyze which products sell together. A bank might analyze transactions to detect fraud. A hospital might analyze appointment patterns to reduce no-shows.

AI is now being used across that pipeline. It helps people clean data, search large datasets, summarize documents, build charts, detect unusual behavior, predict future outcomes, and ask questions in plain English instead of writing database code. The important point is that AI is not magic. It still depends on good data, clear business questions, security rules, and human judgment.

Recent government data gives us a useful reality check. The U.S. Census Bureau added supplemental AI questions to its Business Trends and Outlook Survey from November 2025 through February 2026 to measure how businesses are using AI, including by industry, geography, firm size, task, and impact on work. (Census.gov) A 2026 Census working paper using that survey found that during November 2025 through January 2026, about **18% of firms** used AI in at least one business function, rising to **32% when weighted by employment**, meaning workers are more likely to be employed at firms that use AI than the firm-count number alone suggests. Adoption was much higher in very large firms and in information, professional services, and finance. (Census.gov)

That means AI analytics is real, but uneven. Many large companies are already using it. Many small organizations are experimenting. Plenty of businesses still use ordinary spreadsheets, dashboards, and manual reports.

2.1 Core concepts: AI, machine learning, analytics, and business intelligence

Artificial intelligence, or AI, is software that performs tasks we usually associate with human reasoning, such as recognizing patterns, classifying information, generating text, or making recommendations. **Machine learning**, or ML, is a major type of AI in which a system learns patterns from examples instead of being programmed with every rule manually. For example, a fraud detection model may learn from thousands or millions of past transactions labeled “fraud” or “not fraud.”

Generative AI is AI that creates new content, such as text, code, images, summaries, or answers. A **large language model**, or LLM, is a generative AI model trained to work with language. ChatGPT, Claude, Gemini, and many enterprise assistants are examples of LLM-based systems. In analytics, LLMs are often used as an interface: a manager types, "Which region had the biggest decline in repeat customers?" and the system translates that question into a database query, chart, or written explanation.

Business intelligence, often shortened to BI, refers to software and processes for reporting and dashboards. Power BI, Tableau, Looker, and similar tools help organizations visualize data. A **dashboard** is a screen of charts and metrics, such as daily sales, inventory levels, or customer satisfaction scores. A **data warehouse** is a central database designed for analysis. A **data lake** stores large amounts of raw or semi-structured data. A **semantic model** is a business-friendly layer that defines what terms mean, such as "net revenue," "active customer," or "return rate," so that analytics tools answer questions consistently.

Modern AI analytics usually combines these pieces: databases, dashboards, machine learning models, natural-language chat, data governance, and human review.

2.2 How AI is being used in data mining and analytics

2.2.1 Natural-language analytics: “talking to your data”

One of the most visible changes is natural-language analytics. Instead of writing SQL, which stands for Structured Query Language and is the standard language for asking questions of databases, a user can type a normal question.

For example:

“Show me total profit by product category for the last three months, and highlight anything unusual.”

The AI system may generate a chart, a table, or a written summary. Microsoft says Copilot in Power BI lets users interact with data using natural language, create and analyze visuals, ask ad hoc questions, and summarize reports, while also warning that the data model must be prepared properly or answers can become generic, inaccurate, or misleading. (Microsoft Learn) Microsoft also documents Copilot’s ability to create and edit Power BI report pages from natural-language prompts. (Microsoft Learn)

Google’s BigQuery analytics platform now includes Gemini features that let users use natural language to discover, transform, query, and visualize data. (Google Cloud Documentation) Snowflake’s Cortex Analyst is described as an LLM-powered feature that lets business users ask questions in natural language and receive answers from structured data in Snowflake without writing SQL. (Snowflake Documentation) Tableau Agent, from Salesforce, similarly helps users explore data, create visualizations, and uncover insights through a conversational assistant connected to a workbook or data source. (Tableau Help)

The business value is not just convenience. Natural-language analytics can let nontechnical employees ask follow-up questions without waiting days for an analyst. The risk is that people may trust a fluent answer too quickly. A good analytics team still checks whether the AI used the right table, the right definition, and the right time period.

2.2.2 Data preparation and cleaning

Before data can be analyzed, it usually has to be cleaned. **Data cleaning** means fixing or handling missing values, duplicate records, inconsistent labels, impossible dates, and other quality problems. A simple example: one sales table might say “NY,” another says “New York,” and another says “N.Y.” A person or system has to decide whether these mean the same thing.

AI tools are increasingly used to suggest data transformations. They can detect possible duplicates, standardize categories, recommend joins between tables, and explain confusing fields. Google’s BigQuery release notes reported a 2025 feature allowing users to prepare data with Gemini using natural language. (Google Cloud Documentation) Tableau Agent also advertises natural-language data preparation, including cleaning, shaping, calculations, and pivoting tables. (Tableau)

This is useful because many analytics projects spend more time preparing data than analyzing it. But there is a catch: AI may make a confident-looking cleaning suggestion that changes the meaning of the data. For example, combining "returned item" and "refunded order" might be wrong if the company tracks those differently. In real organizations, data preparation should be reviewed by someone who understands the business.

2.2.3 Text and document mining

Traditional analytics worked best with structured data: rows and columns, like a spreadsheet. But many business records are **unstructured data**, meaning they do not fit neatly into rows and columns. Examples include customer emails, chat transcripts, PDFs, repair notes, contracts, call-center logs, product reviews, and medical notes.

AI is very useful for mining this kind of text. It can classify support tickets, extract names or dates from documents, summarize customer complaints, detect sentiment, and group similar comments. Snowflake's AI SQL functions include tasks such as extracting entities, aggregating insights across customer tickets, sentiment analysis, document parsing, and preparing data for retrieval-augmented generation. (Snowflake Documentation) **Retrieval-augmented generation**, or RAG, is a method where an AI model answers a question using retrieved company documents or database records, rather than relying only on what the model learned during training.

JPMorgan Chase provides a concrete example from a large financial institution. In its 2025 annual report, the company described AI capabilities tying together research, data, and risk across about 90,000 securities and 22 million documents, plus tools that ingest thousands of broker research reports daily and push prioritized insights to investors. These are company-reported claims, not independent academic measurements, but they show how a major bank describes actual AI analytics use in production.

2.2.4 Predictive analytics and forecasting

Predictive analytics means using past data to estimate what may happen next. A model might predict next month's sales, which customers are likely to cancel a subscription, which invoices may be paid late, or which machines may fail soon.

AI forecasting is not new. Businesses have used statistical forecasting for decades. What has changed is that AI tools are becoming easier to use and more embedded in cloud platforms. Google's BigQuery AI documentation says users can train, evaluate, and deploy predictive analytics models directly within BigQuery using SQL, and can use generative AI functions for summarization, sentiment analysis, and enrichment. (Google Cloud)

JPMorgan Chase reported more than ten years of advanced machine learning and AI work, with company-described value across credit, fraud, and personalization. Its 2025 annual report also says its commercial and investment bank uses AI in transaction screening, cash-flow forecasting for treasury clients, securities inventory management, pricing, risk management, and capital optimization. These examples show AI analytics moving beyond "make me a chart" into operational decision support.

2.2.5 Anomaly detection and fraud detection

An **anomaly** is something unusual compared with normal patterns. Anomaly detection is used in cybersecurity, banking, manufacturing, retail, and logistics. A credit card purchase in a new country may be normal for a traveler but suspicious for someone who never leaves town. A sudden spike in product returns may indicate a defective batch or fraud. A server producing unusual login attempts may signal a cyberattack.

Machine learning is useful here because fraud patterns change. Rule-based systems, such as "block every transaction over $1,000," are too simple. AI systems can look at combinations of signals: location, amount, device, time of day, purchase history, merchant type, and recent behavior.

The practical deployment pattern is usually human-in-the-loop. **Human-in-the-loop** means the AI flags or ranks cases, but a person reviews important decisions. This matters because false positives can annoy customers or delay legitimate transactions, while false negatives can miss real fraud.

2.2.6 Customer segmentation, personalization, and recommendations

Segmentation means grouping customers or products based on shared characteristics. A retailer might identify "budget shoppers," "high-frequency snack buyers," or "seasonal gift buyers." A streaming service

might group viewers by viewing patterns. A college might group students by support needs, such as first-generation students who may benefit from early advising.

AI can discover patterns that are hard to spot manually. It can recommend products, personalize offers, rank sales leads, or suggest next-best actions for customer service. Salesforce's 2025 Form 10-K says its platform connects customer data across systems to create a complete view of customers, and its analytics offerings, including Tableau, help users visualize, analyze, spot trends, predict outcomes, receive recommendations, and take action with AI agents. (SEC)

This kind of analytics can be helpful, but it can also become invasive or unfair if organizations use sensitive personal data carelessly. Personalization should not become manipulation, discrimination, or surveillance.

2.2.7 AI agents connected to business data

An **AI agent** is software that can use tools, follow instructions, and take steps toward a goal. In analytics, an agent might answer a question, run a query, create a chart, summarize results, and send a draft report. In more advanced settings, an agent may trigger a workflow, such as opening a support ticket or recommending a reorder quantity.

Snowflake's 2025 Form 10-K says Snowflake Intelligence enables customers to create data agents so business users can take actions on structured and unstructured data without technical knowledge or coding skills. (SEC) Palantir's 2024 Form 10-K says its Artificial Intelligence Platform, or AIP, connects large language models and other AI with customer data and operations to support decision-making within legal, ethical, and security constraints. It also describes workflows involving data engineers, analysts, data scientists, business users, and senior leaders. (SEC)

This is where analytics starts to blend with operations. The system is not only answering "What happened?" It is helping decide "What should we do next?"

2.3 What deployment looks like inside a real organization

A typical business AI analytics setup has several layers.

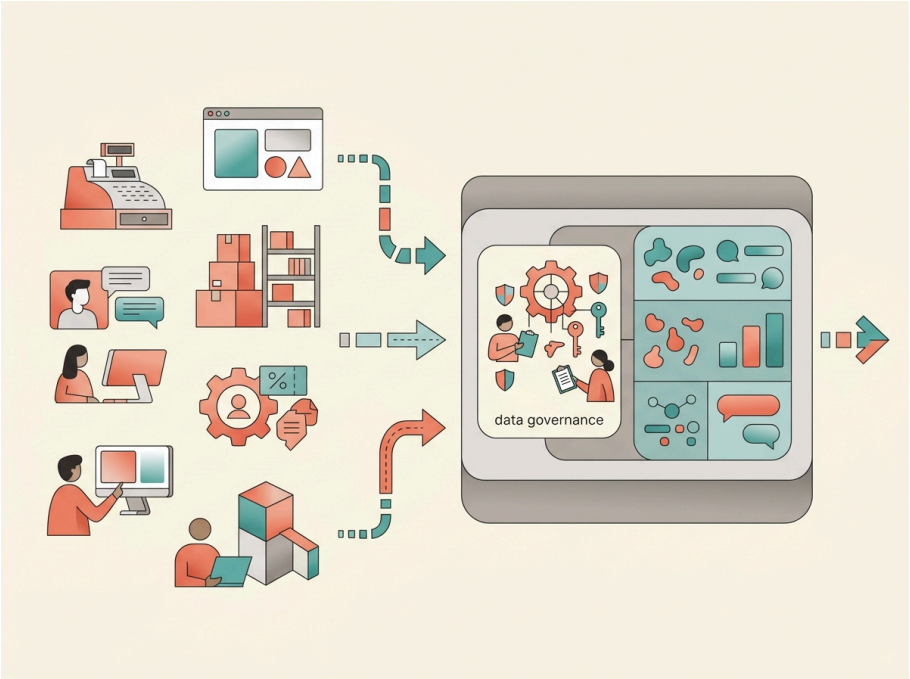

First, data is collected from business systems: sales systems, websites, inventory systems, customer relationship management systems, accounting software, support tickets, sensors, or external market data. Second, the data is moved into a warehouse, lake, or platform where it can be governed. **Data governance** means rules for who can access data, what definitions mean, how quality is checked, and how data is protected. Third, analytics tools create reports, dashboards, models, and alerts. Fourth, AI features help users ask questions, summarize results, detect patterns, and recommend actions.

This is why companies selling analytics platforms now describe themselves as AI data platforms. Alphabet's 2025 10-K says Google Cloud's data and analytics offerings help customers migrate, clean,

prepare, and feed data into models, while unifying data lakes, warehouses, governance, and machine learning so users can analyze data using AI models across clouds. (SEC) Snowflake says its platform helps customers consolidate data into a single source of truth, drive insights, apply AI to business problems, build data applications, and handle structured, semi-structured, and unstructured data. (SEC)

On the user-company side, Walmart's 2025 Form 10-K says it continues to invest in AI and generative AI to enhance customer shopping and associate work experiences and improve efficiencies in supply chain, operations, management functions, and talent recruitment and development. Walmart also warns that these technologies bring operational, legal, privacy, compliance, and reputation risks. (Walmart Inc.) That is a realistic disclosure: AI is useful, but it creates new responsibilities.

2.4 What AI does well — and what it does poorly

AI analytics is strong at pattern recognition, summarization, classification, prediction, and interface design. It can speed up routine tasks and make data tools more accessible.

But there are major limits.

Garbage in, garbage out means bad input data leads to bad output. If a company's customer records are outdated, duplicated, or biased, AI may produce polished but wrong conclusions. **Hallucination** means a generative AI system produces an answer that sounds confident but is false or unsupported. In analytics, hallucination is especially dangerous because a false number in a business report can lead to a bad decision.

Bias means a system produces unfair results for certain groups, often because the training data reflects past discrimination or incomplete measurement. **Model drift** means a model becomes less accurate over time because the world changes. For example, a demand forecast trained on last year's shopping patterns may fail after a competitor opens nearby or prices change.

Regulators and standards bodies are paying attention. NIST's Generative AI Profile, released in 2024, is designed to help organizations identify generative AI risks and choose risk-management actions. (nist.gov) The European Union's AI Act entered into force in 2024 and becomes broadly applicable in stages, with full applicability generally beginning in August 2026, though some provisions apply earlier. (Digital Strategy EU) The U.S. Federal Trade Commission has also warned that companies may face consequences for using consumer data unlawfully, including deletion of products, models, or algorithms built from unlawfully obtained data. (Federal Trade Commission)

The practical lesson is simple: analytics AI should be checked, logged, secured, and governed.

2.5 Hands-on lab: Use AI to mine a small sales dataset

2.5.1 Lab goal

You will use an AI assistant and a spreadsheet to analyze a small dataset. The goal is not to let AI do everything. The goal is to learn how to ask better questions, verify answers, and turn data into a short business recommendation.

You may use Excel, Google Sheets, LibreOffice Calc, or another spreadsheet tool. You may use any AI assistant your class allows.

2.5.2 Scenario

You are helping a small campus store analyze product sales. The manager wants to know which products are profitable, which have high return rates, and what actions to take next month.

Copy this CSV data into a spreadsheet:

```
month,channel,category,product,units_sold,revenue,cost,returns,discount_pct
2026-01,Store,Beverages,Cold Brew,520,2080,780,12,0.05
2026-01,Online,Beverages,Cold Brew,180,756,270,7,0.10
2026-01,Store,Snacks,Protein Bar,430,1075,516,9,0.00
2026-01,Online,Snacks,Protein Bar,210,546,252,13,0.05
2026-01,Store,Supplies,Notebook,260,1040,572,5,0.00
2026-01,Online,Supplies,Notebook,160,704,352,11,0.15
2026-02,Store,Beverages,Cold Brew,610,2440,915,14,0.05
2026-02,Online,Beverages,Cold Brew,240,1008,360,8,0.10
2026-02,Store,Snacks,Protein Bar,390,975,468,10,0.00
2026-02,Online,Snacks,Protein Bar,260,676,312,19,0.05
```

2026-02,Store,Supplies,Notebook,300,1200,660,6,0.00
2026-02,Online,Supplies,Notebook,190,836,418,16,0.15
2026-03,Store,Beverages,Cold Brew,700,2800,1050,18,0.05
2026-03,Online,Beverages,Cold Brew,310,1302,465,10,0.10
2026-03,Store,Snacks,Protein Bar,420,1050,504,8,0.00
2026-03,Online,Snacks,Protein Bar,300,780,360,24,0.05
2026-03,Store,Supplies,Notebook,340,1360,748,7,0.00
2026-03,Online,Supplies,Notebook,230,1012,506,21,0.15

2.5.3 Step 1: Create calculated fields

Add these columns:
profit = revenue - cost
profit_margin = profit / revenue
return_rate = returns / units_sold
net_units = units_sold - returns
Use spreadsheet formulas. Format profit_margin and return_rate as percentages.

2.5.4 Step 2: Ask AI for a data profile

Paste the dataset into your AI assistant and use this prompt:
You are helping me analyze a small campus store dataset.
First, profile the data. Identify the columns, explain what each one means, and suggest three business questions this dataset can answer.
Do not make conclusions yet.
Check whether the AI correctly understands the dataset. If it invents columns or makes claims that are not in the data, note that as an error.

2.5.5 Step 3: Ask AI for analysis, but require evidence

Use this prompt:
Using only the data I provided, answer these questions:
1. Which product-category combination produced the most total profit?
2. Which channel has the higher average return rate?
3. Which product seems most concerning from a returns perspective?
Show the calculations or explain exactly how you reached each answer.
Now verify the answer in your spreadsheet using a pivot table or manual sorting. This is the key skill: do not just accept the AI's answer because it sounds professional.

2.5.6 Step 4: Ask for a business recommendation

Use this prompt:
Write a short business recommendation for the campus store manager.
Include:
- one action to increase profit,
- one action to reduce returns,
- one warning about the limits of this dataset.
Keep it under 180 words and avoid unsupported claims.
A strong answer should mention that the dataset is small, synthetic, and only covers three months. It should avoid pretending to know customer demographics, weather, supplier problems, or long-term demand.

2.5.7 Step 5: Write your lab memo

Submit a one-page memo with four parts:

1. **Pattern found:** What did the data show?
2. **AI help:** What did the AI assistant do well?
3. **AI risk:** What did the AI get wrong, skip, or overstate?

4. **Decision:** What would you recommend to the manager?

This lab mirrors real AI analytics work. The AI can speed up exploration, but the analyst remains responsible for checking the logic.

2.6 Key takeaways

AI tools are definitely being used in data mining and analytics. The strongest evidence is not futuristic advertising; it is visible in government adoption surveys, public company filings, and product documentation from major data platforms.

The most common uses are natural-language data questions, automated reporting, data cleaning, text mining, forecasting, anomaly detection, recommendations, and AI agents connected to business data. The pattern across tools is similar: AI is becoming part of the analytics workflow rather than a separate "robot analyst."

The realistic picture is mixed. AI can reduce friction and help more people use data. It can also produce wrong answers faster, hide bias behind polished language, or expose sensitive information if governance is weak. Good organizations treat AI analytics as a powerful assistant inside a controlled data process, not as an unquestionable decision-maker.

3. AI Tools in Digital Marketing

Course Outcome

VCCS-3. Exploring Marketing & Personalization Using AI-tools: Evaluate the impact of AI on marketing and personalization, including content customization, audience targeting, and campaign optimization.

AI tools are already being used heavily in digital marketing. They are not only "future technology," and they are not limited to giant technology companies. In 2026, AI shows up inside the everyday tools marketers use to create ads, write email campaigns, choose audiences, set bids, test creative versions, personalize websites, analyze customer behavior, and summarize campaign results.

The important thing to understand is that most companies are **not** building their own large AI systems from scratch. They are usually using AI that is already built into platforms such as Google Ads, Meta Ads, Amazon Ads, Adobe, HubSpot, Mailchimp, ecommerce systems, customer relationship management software, and analytics tools. In other words, businesses often "rent" AI through the marketing software they already pay for.

Recent evidence shows this is not a small trend. The February 2025 CMO Survey reported that artificial intelligence and machine learning powered 17.2% of marketing efforts among surveyed U.S. companies, roughly double the 2022 level, and marketing leaders expected that share to keep growing. The 2026 CMO Survey said marketers expected AI to account for more than half of marketing activities within three years, while also warning that many organizations were not fully ready to manage the change. These are survey findings, not audited performance data, but they show how mainstream AI has become in marketing leadership conversations. (Duke's Fuqua School of Business)

At the same time, the advertising industry's own research shows a more cautious reality. A 2025 Interactive Advertising Bureau study found that only 30% of agencies, brands, and publishers had fully integrated AI across the media campaign life cycle, and many still lacked a strategic roadmap. Another IAB study found that more than half of marketers were using generative AI for creative content and audience targeting, but more than 70% had already encountered an AI-related incident such as inaccurate content, bias, or off-brand material. (IAB)

So the best answer is: **AI is widely used in digital marketing, but it is usually used as a powerful assistant and automation layer, not as a fully independent marketing department.**

3.1 What “digital marketing” means

Digital marketing means promoting products, services, brands, or ideas through digital channels. These channels include search engines, websites, social media, email, mobile apps, online marketplaces, streaming video, digital audio, and connected TV. A connected TV is a television that streams internet-based video through apps such as YouTube, Hulu, Netflix, or other streaming services.

Digital marketing usually follows a basic path called the **marketing funnel**. A funnel is a simplified model of how people move from not knowing about a product to possibly becoming customers. The main stages are:

1. **Awareness** — people first learn that a product or brand exists.
2. **Consideration** — people compare options and decide whether they are interested.
3. **Conversion** — people take the desired action, such as buying, signing up, downloading, donating, or booking.
4. **Retention** — people return, renew, subscribe, or buy again.

A **conversion** is the action a business wants a person to take. For an online store, a conversion might be a purchase. For a dentist’s office, it might be booking an appointment. For a college, it might be requesting information or starting an application.

AI can help at every stage of this funnel. It can generate awareness ads, recommend which audiences should see them, personalize offers, decide when to send emails, predict which customers might leave, and summarize which parts of a campaign worked.

3.2 What “AI tool” means in marketing

Artificial intelligence, or **AI**, refers to computer systems that perform tasks associated with human judgment, pattern recognition, language, prediction, or decision support. In digital marketing, three kinds of AI matter most.

Machine learning is a type of AI where software learns patterns from data rather than being programmed with every rule by hand. For example, an ad platform can learn that people who recently searched for “running shoes for flat feet” are more likely to click on a shoe ad than people browsing unrelated content.

Predictive AI estimates what is likely to happen. It might predict which customers are likely to buy, which subscribers may cancel, which ad placement may get clicks, or how much a company should bid for an ad impression.

Generative AI creates new content, such as text, images, video, audio, summaries, product descriptions, or ad variations. A marketer might ask a generative AI tool to write five email subject lines, create a product image background, or draft a 15-second video script.

A **prompt** is the instruction a person gives to an AI system. For example: “Write three Instagram captions for a local bakery promoting a weekend cupcake sale. Keep the tone friendly and under 100 characters.”

A **hallucination** is when an AI system produces information that sounds confident but is false, unsupported, or made up. In marketing, hallucinations are especially risky because an AI tool might invent a discount, misstate a product feature, create a fake quote, or make a legal claim the company cannot support.

3.3 Where AI is actually being used

3.3.1 Ad targeting and automated bidding

One of the most important uses of AI in digital marketing is deciding **who sees which ad, when, and at what price**.

In online advertising, businesses often buy ad space through automated systems. This is called **programmatic advertising**, meaning ad placements are bought and sold using software, often through real-time auctions. When a person opens a website, app, search results page, or video platform, ad systems may instantly decide which advertiser’s message should appear.

AI helps with this process by analyzing signals such as search terms, location, device type, browsing behavior, past purchases, content context, and advertiser goals. The AI does not “know” the person like a friend would. Instead, it estimates probabilities: Who is likely to click? Who is likely to buy? Which ad version is likely to perform best?

Google’s Performance Max campaigns are a clear example. Google describes Performance Max as using Google AI across bidding, budget optimization, audiences, creative, attribution, and other parts of campaign management. The system can optimize across Google properties based on goals such as conversions, cost per acquisition, or return on ad spend. (Google Help)

Cost per acquisition, or **CPA**, means the average amount spent to get one conversion. If a company spends $200 and gets 10 purchases, the CPA is $20. **Return on ad spend**, or **ROAS**, compares revenue to advertising cost. If a company spends $100 on ads and earns $400 in revenue from those ads, the ROAS is 4:1.

Meta also emphasizes AI as central to its advertising business. In its 2024 annual filing, Meta said AI investments power ranking systems, its discovery engine, and tools advertisers use to reach customers. Meta also said it generates substantially all of its revenue from selling ads to marketers. (SEC)

This matters because Meta, Google, Amazon, TikTok, and other large platforms have enormous incentives to automate advertising. The easier they make campaign setup, targeting, creative production, and optimization, the more businesses can spend money on ads without needing a large internal marketing team.

3.3.2 Creative production: text, images, and video

The most visible use of AI in marketing is creative production. **Creative** means the words, images, videos, layouts, and messages people see in an ad or campaign.

Before generative AI, producing ad creative often required several separate steps: write a brief, hire designers or video editors, schedule a photo shoot, edit assets, create different versions for each platform, send files for approval, and upload the finished work. That process still exists, especially for major brand campaigns, but AI has changed the speed and scale of everyday production.

Generative AI can create:

- ad headlines;
- email subject lines;
- product descriptions;
- social media captions;
- image backgrounds;
- video storyboards;
- short product videos;
- translated campaign copy;
- many versions of the same message for different audiences.

This is not theoretical. Reuters reported in 2024 that Klarna used generative AI for marketing campaigns and image generation, estimating about $10 million in annual savings. Klarna said AI helped reduce image production costs by about $6 million, shortened image development cycles from six weeks to seven days, and helped produce more than 1,000 images in three months. (Reuters)

A 2026 Reuters report described Polish fashion retailer LPP using AI to predict trends and generate marketing visuals. According to the company, AI-generated visuals rose from 20% of marketing visuals in 2025 to 80%, while content costs fell by 60%. LPP also said AI helped shorten parts of the design process from months to weeks. (Reuters)

Amazon Ads has also moved directly into AI-generated creative. In 2025, Amazon made an enhanced Video Generator available to U.S. advertisers. The tool can turn a product image into short video options with motion, text animations, music, and brand elements. Amazon describes these tools as reducing "creative friction," meaning they make it easier for advertisers to create usable assets without a full production team. (Amazon Ads)

Adobe is another important example because many professional marketing teams already use Adobe software. Adobe's 2025 annual report describes Firefly as a generative AI system for images, video, vector graphics, audio, design templates, and other assets. Adobe also describes GenStudio as a system for managing the content supply chain, including ideation, production, activation, asset management, and analytics.

A **content supply chain** is the full process of planning, creating, approving, storing, distributing, and measuring marketing content. Large companies may need thousands of content variations: different languages, regions, product lines, platforms, audience segments, and legal requirements. AI is attractive because it can produce and adapt versions faster than traditional workflows.

But faster does not automatically mean better. AI-generated creative can be bland, inaccurate, legally risky, or off-brand. AdExchanger reported in 2025 that Amazon's AI creative tools could still produce outputs that looked "off," and Amazon said only about half of generated assets were actually used. That is a useful reality check: many AI outputs are drafts, not finished professional work. (AdExchanger)

3.3.3 Personalization in email, websites, and customer journeys

Another major use of AI is **personalization**, which means changing a message, offer, recommendation, or experience based on what is known about a customer or visitor.

Personalization can be simple, such as putting a first name in an email. AI-enabled personalization is more advanced. It might decide:

- which products to recommend;
- which email subject line to send;
- when to send a message;
- whether a customer should receive a discount;
- which website banner to show;
- which customers may stop buying soon;
- which customers should be excluded from a campaign because they already purchased.

A **customer relationship management system**, or **CRM**, is software that stores information about customers, prospects, sales activity, service history, and marketing interactions. HubSpot's 2024 annual report describes its Breeze AI as part of a customer platform that includes CRM, marketing automation, engagement hubs, AI agents, and data enrichment. (SEC)

Marketing automation means using software to send messages or trigger actions based on rules or customer behavior. For example, an online store might automatically send a welcome email after signup, a reminder after someone abandons a shopping cart, and a thank-you message after purchase.

Mailchimp, a widely used email marketing platform, offers AI-assisted flow templates that can create designed, on-brand emails for certain users. Its documentation says the AI can use brand kit information and generate content and style suggestions, while users review and activate the automation. (Mailchimp)

This last point is important: the human marketer still matters. A responsible workflow does not allow the AI to invent offers, send messages to the wrong audience, or ignore unsubscribe rules. Humans define the goal, review the content, check the data, and monitor the results.

3.3.4 Search marketing and AI search results

Search marketing is also changing. **Search engine optimization**, or **SEO**, means improving website content so it can be found through search engines. **Search advertising** means paying to show ads when people search for certain topics, products, or questions.

AI affects both.

On the advertising side, Google has introduced AI Max for Search campaigns. Google describes AI Max as using landing page content, existing ads, assets, and generative AI to create customized ad copy that better matches user searches. (Google Help)

On the organic search side, AI-generated summaries can change how people behave after searching. Pew Research Center analyzed browsing behavior from 900 U.S. adults in March 2025 and found that users who saw a Google AI summary clicked a traditional search result in 8% of visits, compared with 15% of visits without an AI summary. Pew also found users clicked links inside the AI summaries only 1% of the time and were more likely to end the browsing session after seeing a summary. (Pew Research Center)

For marketers, that means website traffic from search may become harder to predict. A person might get an answer directly from an AI summary instead of clicking to the company's website. This creates pressure for marketers to write clearer pages, maintain accurate product data, use structured information, and think about visibility inside AI-generated answers, not only traditional search rankings.

Google has also been expanding advertising into AI-powered search experiences. At Google Marketing Live 2025, Google said it was expanding ads in AI Overviews to desktop and bringing ads to AI Mode. (blog.google)

3.3.5 Analytics, attribution, and campaign reporting

AI is also used after campaigns launch. Marketers need to know what worked, what failed, and what to change.

A **key performance indicator**, or **KPI**, is a metric used to judge performance. Common digital marketing KPIs include impressions, clicks, click-through rate, conversion rate, cost per conversion, revenue, and return on ad spend.

An **impression** means an ad or piece of content was shown. A **click-through rate**, or **CTR**, is the percentage of impressions that turned into clicks. A **conversion rate** is the percentage of visitors or clicks that turned into conversions.

A simple version looks like this:

Click-through rate = clicks ÷ impressions
Plain-English translation: "Out of everyone who saw this, what share clicked?"
Conversion rate = conversions ÷ clicks
Plain-English translation: "Out of everyone who clicked, what share took the desired action?"

Attribution means assigning credit for a conversion to one or more marketing touchpoints. For example, suppose someone sees a TikTok ad, later searches the brand on Google, clicks an email, and finally buys. Which channel gets credit? The answer is not always obvious.

AI tools can help summarize campaign data, detect unusual changes, forecast performance, group customers into segments, and recommend budget shifts. But attribution is one of the easiest places to fool yourself. A platform may report that its own ads drove sales, but some customers might have purchased anyway. That is why serious marketers use controlled tests, holdout groups, incrementality studies, or careful before-and-after comparisons whenever possible.

Incrementality means the extra result caused by marketing that would not have happened otherwise. If 100 people buy after seeing an ad, but 60 of them would have bought anyway, the incremental effect is closer to 40 purchases, not 100.

3.4 How this looks inside a business

A practical AI marketing setup usually has five layers.

First is **data**. This includes product information, website visits, email engagement, purchase history, search terms, ad performance, loyalty data, and customer service records. **First-party data** is data a business collects directly from its own customers or users, such as email signups, purchase records, or app activity.

Second is the **tool layer**. This may include ad platforms, email platforms, CRM systems, ecommerce software, analytics dashboards, creative tools, and generative AI assistants.

Third is **activation**. Activation means putting the campaign into the world: launching ads, sending emails, updating website content, publishing social posts, or changing product recommendations.

Fourth is **measurement**. The business tracks what happened: impressions, clicks, purchases, unsubscribes, revenue, complaints, or customer lifetime value. **Customer lifetime value** is an estimate of how much revenue or profit a customer may generate over the full relationship with the business.

Fifth is **governance**. Governance means the rules, approvals, roles, and monitoring that keep AI use accurate, legal, ethical, and brand-safe. **Brand safety** means avoiding situations where ads or content appear in harmful, misleading, offensive, or inappropriate contexts.

The companies that use AI well do not simply say, "Let the AI handle marketing." They build workflows. They decide who can generate content, who approves it, what data may be used, what claims require legal review, and what happens if the AI output is wrong.

3.5 How AI use differs by business size

Small businesses often use AI through built-in platform features. A local restaurant, boutique, repair shop, or fitness studio may use AI inside Google Ads, Meta Ads, Mailchimp, Shopify, Square, Canva, or other tools. The business owner may never train a model or write code. They might use AI to draft posts, create ad variations, design email promotions, or let an ad platform optimize for calls, reservations, or purchases.

Medium-sized businesses usually connect more systems. They may use a CRM, ecommerce database, email automation platform, ad platforms, website analytics, and customer support tools. Their challenge is less about generating one ad and more about keeping customer data clean, avoiding duplicate messages, and measuring what actually caused sales.

Large enterprises often focus on scale, compliance, and brand control. A global consumer brand may need product images in dozens of countries, multiple languages, different legal requirements, and different retail partners. That is why companies such as Adobe emphasize content supply chain tools and brand-controlled generative AI systems.

Large companies are also using AI to reduce production costs and speed up campaign cycles. Unilever, in company-published materials, described using AI and digital twins of products to create marketing imagery faster and cheaper. Because this is company-published, it should be treated as the company's own account, not independent proof of performance. (Unilever)

A **digital twin** is a digital representation of a real-world object, process, or system. In marketing, a digital twin of a product might allow a team to generate many realistic product images without photographing the product every time.

3.6 Risks and limits

AI marketing has real benefits, but it also creates real risks.

The first risk is accuracy. Generative AI can invent product features, fake statistics, false claims, or misleading comparisons. In regulated industries such as finance, healthcare, insurance, education, and housing, inaccurate marketing claims can create legal problems.

The second risk is fake social proof. Social proof means evidence that other people like or trust something, such as reviews, testimonials, follower counts, or ratings. In 2024, the U.S. Federal Trade Commission finalized a rule banning fake reviews and testimonials, including reviews that misrepresent a nonexistent person or someone's actual experience. The rule also covers buying, selling, or spreading fake reviews when the business knew or should have known they were false. (Federal Trade Commission)

The third risk is deceptive AI claims. The FTC's 2024 "Operation AI Comply" actions targeted companies accused of using AI claims in misleading or unfair ways, including fake review tools and exaggerated AI-powered business opportunities. (Federal Trade Commission)

The fourth risk is privacy and discrimination. AI targeting can become harmful if it uses sensitive traits, excludes protected groups unfairly, or relies on data collected without meaningful consent. In the European Union, the Digital Services Act requires online ads to be clearly labeled, requires very large online platforms to maintain public ad repositories, and restricts ad targeting based on sensitive data such as race, religion, or sexual orientation. (Digital Strategy EU)

The fifth risk is overreliance on platform numbers. Google, Meta, Amazon, and other platforms provide useful reporting, but they are also selling the ads. A smart business treats platform dashboards as evidence,

not as the final truth. Whenever possible, marketers compare platform data with sales records, website analytics, customer surveys, experiments, and financial results.

The sixth risk is workflow disruption. When AI can create many images, draft many messages, or automate media buying, some tasks become faster and cheaper. That can help small teams, but it can also reduce work for freelancers, agencies, translators, photographers, junior copywriters, and production vendors. Klarna's reported savings from reducing external marketing suppliers show that AI adoption is not only a technical change; it is also an organizational and labor-market change. (Reuters)

3.7 Hands-on lab: Build an AI-assisted digital marketing campaign

3.7.1 Lab goal

In this lab, you will design a small AI-assisted digital marketing campaign for a fictional local business. You will not need paid advertising accounts or real customer data. The goal is to practice how marketers use AI responsibly: as a research, drafting, testing, and planning assistant.

3.7.2 Scenario

You are helping **Campus Cup**, a small coffee cart near a community college library. Campus Cup wants to promote a new product: an iced latte flight with three small flavors. The product costs $6.50. It is available Monday through Thursday from 8:00 a.m. to 11:00 a.m. The business wants students, staff, and commuters to preorder online and pick up at the cart.

Campaign goal: Get 120 preorders in two weeks.

Budget: $150.

Channels: Instagram-style social posts, search ads, and one email campaign.

Important constraints: No delivery. No medical claims. No fake reviews. No "best coffee on campus" claim unless the business can prove it. Mention that oat milk is available, but do not claim the drink is allergen-free.

3.7.3 Step 1: Write the campaign brief

A **campaign brief** is a short document that explains the goal, audience, message, channels, and measurement plan.

Create a one-page brief with these sections:

- Product
- Audience
- Main customer problem
- Main promise
- Offer
- Channels
- Budget
- Success metric
- Risks to avoid

Use this prompt with an AI writing tool:

You are helping a small coffee cart create a two-week digital marketing campaign.
Product: iced latte flight with three small flavors.
Price: $6.50.
Availability: Monday through Thursday, 8:00 a.m. to 11:00 a.m.
Location: near a community college library.
Goal: 120 online preorders in two weeks.
Budget: $150.
Audience: students, staff, and commuters.
Constraints: no delivery, no fake reviews, no unsupported "best coffee" claim, oat milk available but not allergen-free.
Create a simple campaign brief for an introductory marketing class.

Use clear language and include risks to avoid.

After the AI responds, revise the brief yourself. Remove anything that sounds exaggerated, unsupported, or confusing.

3.7.4 Step 2: Create audience segments

An **audience segment** is a group of people with similar needs, behaviors, or situations. Ask the AI to suggest three segments.

Prompt:

Based on the Campus Cup campaign brief, suggest three audience segments.

For each segment, include:

1. What they care about
2. What might stop them from buying
3. A message angle that could appeal to them

Keep the answer realistic and avoid stereotypes.

A strong answer might include commuters who want speed, students who want a small treat before class, and staff who want a convenient morning break. A weak answer would rely on stereotypes or make claims without evidence.

3.7.5 Step 3: Draft ad variations

Now create different ad versions. This is one of the most common real uses of generative AI in marketing: producing many first drafts quickly.

Prompt:

Write ad copy for the Campus Cup iced latte flight campaign.

Create:

- 4 short social media captions under 100 characters
- 4 search ad headlines under 35 characters
- 2 search ad descriptions under 90 characters
- 3 calls to action

Rules:

Do not claim “best coffee.”

Do not invent reviews.

Do not imply delivery.

Mention preorder and library pickup when useful.

Use a friendly but not childish tone.

Review the output. Choose the best versions and edit them. Good marketing copy is clear, truthful, and specific. The AI may give you usable drafts, but you are responsible for the final message.

3.7.6 Step 4: Plan one A/B test

An **A/B test** compares two versions of something to see which performs better. The key rule is to test one major difference at a time.

Create two social captions:

- Version A emphasizes convenience.
- Version B emphasizes trying three flavors.

Then define the success metric. For this campaign, the best metric is not likes. It is preorders.

Use this simple structure:

A/B test question:

Does a convenience message or a flavor-variety message get more preorder clicks?

Version A:

[caption]

Version B:

[caption]

Success metric:

Preorder conversion rate.

Decision rule:

After both versions receive a similar number of impressions, choose the one with the higher preorder conversion rate.

3.7.7 Step 5: Build the email

Email marketing is still a major digital marketing channel because the business owns the customer relationship more directly than it does on social platforms.

Prompt:
Write a short promotional email for Campus Cup.
Audience: students and staff who opted in to receive emails.
Goal: get preorders for the iced latte flight.
Include:
- subject line
- preview text
- body copy under 150 words
- one call to action
Rules:
No fake urgency unless tied to the real two-week campaign.
No unsupported claims.
Mention pickup near the library.
Now revise the email. Make sure it is easy to scan on a phone. Make sure the call to action is obvious.

3.7.8 Step 6: Create a measurement plan

Use these metrics:

Impressions — how many times the ad or post was shown. **Clicks** — how many people clicked. **CTR** — clicks divided by impressions. **Conversions** — completed preorders. **Conversion rate** — conversions divided by clicks. **CPA** — ad spend divided by conversions. **Revenue** — conversions multiplied by $6.50. **ROAS** — revenue divided by ad spend.

For example, if the campaign spends $150 and gets 120 preorders:
Revenue = 120 × $6.50 = $780 ROAS = $780 ÷ $150 = 5.2

Plain-English translation: for every $1 spent on ads, the campaign produced $5.20 in preorder revenue before considering costs such as ingredients, labor, payment fees, and time.

3.7.9 Step 7: Run a human review checklist

Before launching any AI-assisted campaign, complete this checklist:

- Are all claims true?
- Did the AI invent testimonials, ratings, awards, or discounts?
- Is the offer clear?
- Are the dates and times accurate?
- Are accessibility needs considered, such as readable text and alt text for images?
- Does the campaign avoid targeting sensitive personal traits?
- Does the email go only to people who opted in?
- Is there an unsubscribe option?
- Would the business be comfortable explaining how the campaign was made?

This is the difference between using AI as a tool and letting AI become a liability.

3.7.10 Lab deliverables

Submit:

1. One-page campaign brief
2. Three audience segments
3. Four social captions

4. Four search ad headlines
5. One promotional email
6. One A/B test plan
7. One measurement plan
8. One human review checklist with at least three revisions you made after AI output

3.8 Chapter takeaway

AI tools are absolutely being used in digital marketing, but the real story is more practical than futuristic. Businesses use AI to create more content versions, automate ad targeting and bidding, personalize emails and websites, summarize analytics, and speed up campaign workflows. Small businesses usually access these capabilities through built-in tools. Large enterprises use AI across content supply chains, customer data systems, ad platforms, and approval workflows.

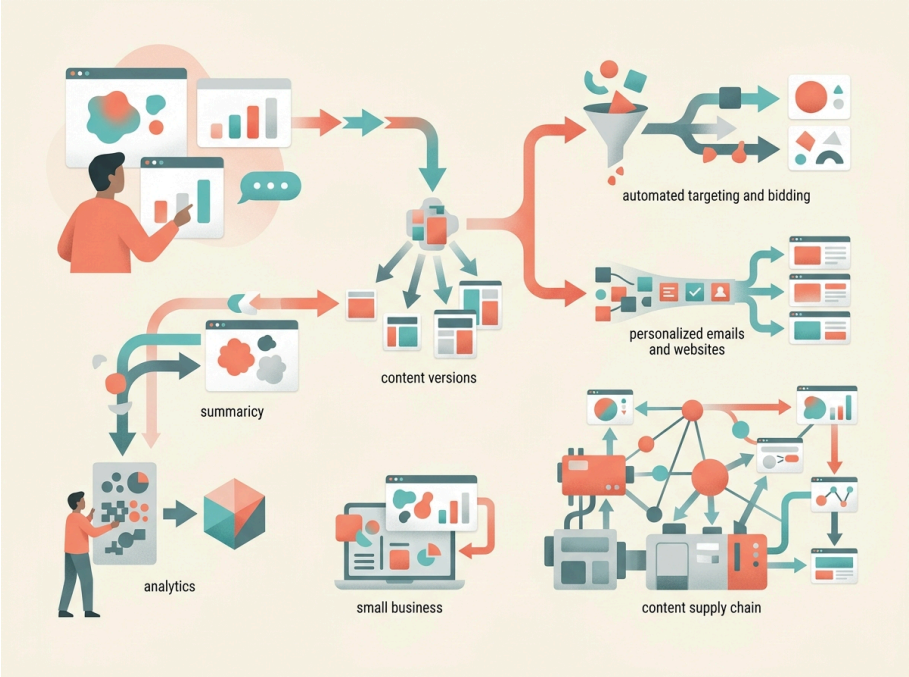

The best marketers in 2026 are not simply asking AI to "make ads." They are learning how to give better instructions, use cleaner data, test results carefully, protect customer privacy, avoid fake or misleading content, and keep humans responsible for judgment. AI can make digital marketing faster and more scalable. It does not remove the need for strategy, ethics, creativity, or accountability.

4. AI Tools in Fraud Detection and Cybersecurity

Course Outcome

VCCS-4. Explain how AI enhances fraud detection and security by identifying fraudulent transactions, cybersecurity threats, and suspicious activities.

AI tools are absolutely being used in fraud detection and security. In fact, these are among the most common real-world uses of artificial intelligence in business today. But they usually do not look like a robot "catching criminals." They look like risk scores, alerts, dashboards, automated blocks, identity checks, phishing triage, and investigation summaries that help human workers decide what to do next.

A useful way to think about AI in fraud and security is this: modern businesses see too many transactions, logins, emails, files, devices, and network events for people to inspect one by one. AI tools help sort that flood of activity into "probably normal," "possibly risky," and "urgent." The U.S. Government Accountability Office, for example, reports that financial institutions use AI for activities such as detecting illicit activity, identifying fake IDs, screening against sanctions lists, analyzing transaction data, and detecting cyber threats in real time. (GAO Files)

Fraud detection means finding deception that could lead to stolen money, stolen goods, stolen accounts, or illegal transactions. Security, or cybersecurity, means protecting systems, devices, networks, and data from unauthorized access or harm. These two areas overlap. A stolen password can become credit card fraud. A phishing email can lead to ransomware. A fake identity can be used to open a bank account, launder money, or attack a business.

The scale of the problem explains why AI is so attractive. The Federal Trade Commission reported that U.S. consumers lost more than $12.5 billion to fraud in 2024, with investment scams and imposter scams among the largest categories. (Federal Trade Commission) The FBI's Internet Crime Complaint Center reported more than 1 million complaints in 2025 and nearly $21 billion in losses from cyber-enabled crime; it also created a specific AI-related section after receiving more than 22,000 AI-related complaints with nearly $893 million in reported losses. (Federal Bureau of Investigation)

So the basic answer is: yes, AI is being used, but mostly as a high-speed decision-support layer. It helps businesses notice patterns faster than people can, but it still needs good data, sensible rules, human oversight, and careful governance.

4.1 The basic idea: AI turns messy activity into risk signals

Artificial intelligence, or AI, is software designed to perform tasks that normally require human-like judgment, such as recognizing patterns, classifying information, or generating explanations. Machine learning is a type of AI in which software learns patterns from data instead of being programmed with only fixed rules. A model is the trained system that makes predictions or classifications.

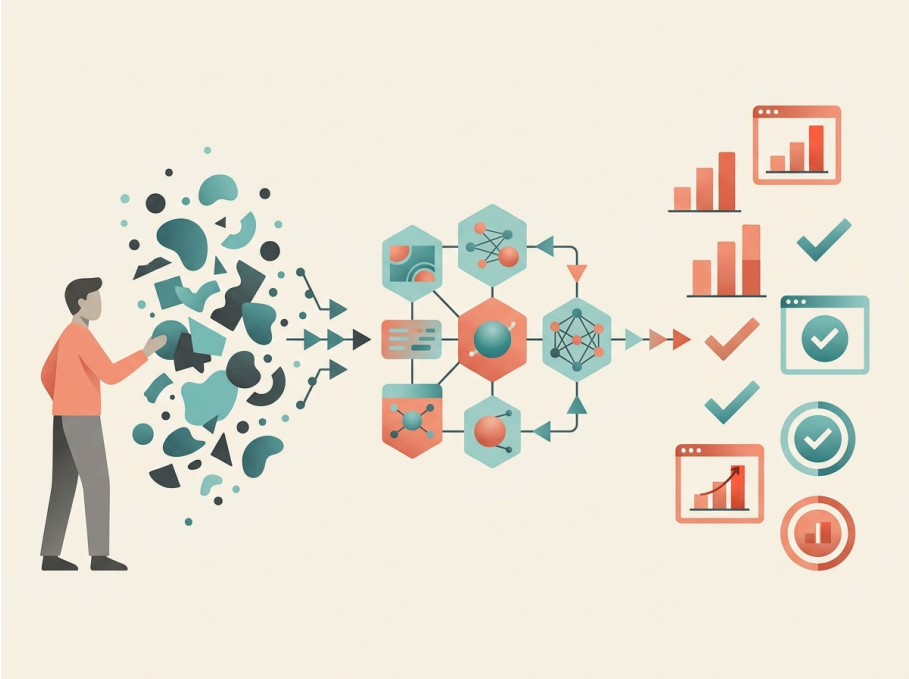

In fraud and security, the model usually works with features. A feature is a piece of information the system can use, such as transaction amount, time of day, location, device type, number of failed login attempts, or whether a shipping address matches a billing address.

For example, imagine an online order. A fraud system might consider these signals:

The customer is using a new device. The order is unusually expensive. The shipping address is new. The account was created yesterday. The customer tried three credit cards before one worked. The IP address appears to come from a country where the customer has never logged in before.

None of these facts alone proves fraud. A real customer might buy a laptop while traveling. But the combination may be suspicious. AI is good at combining many weak clues into a risk score. A risk score is a number that estimates how risky an event looks. A payment might receive a score of 92 out of 100, meaning "very suspicious," or 12 out of 100, meaning "probably normal."

That score can trigger different actions. A low-risk transaction may be approved automatically. A medium-risk transaction may require step-up authentication, meaning the customer must provide extra proof, such as a passkey, one-time code, or app confirmation. A high-risk transaction may be declined or sent to a human fraud analyst.

This is why AI in fraud detection is rarely a single yes-or-no machine. It is usually part of a workflow: collect signals, score the event, apply business rules, send difficult cases to people, and learn from the final outcome.

4.2 Why old-fashioned rules are not enough

Businesses still use rules. A rule is a fixed instruction such as "block any order above $5,000 from a brand-new account" or "require extra verification after five failed login attempts." Rules are useful because they are easy to understand and quick to apply.

But fraudsters adapt. Once criminals learn a rule, they change their behavior. If orders above $5,000 are blocked, they may place five orders for $900. If a stolen card is tested with large purchases, they may first test it with a $2 charge. If an email filter blocks one phrase, attackers write a different phrase.

Machine learning helps because it can detect patterns across many features at once. It can learn that a transaction is suspicious not because of one obvious rule, but because it resembles past fraud cases in subtle ways. It can also be updated as criminals change tactics.

That said, AI does not eliminate rules. In real systems, rules and AI are usually combined. A bank may use machine learning to assign a risk score, then use rules to decide what happens at each score level. For example, scores under 30 might pass, scores from 30 to 70 might require extra verification, and scores above 70 might go to manual review.

4.3 How AI is used in payment fraud

Payment fraud is one of the clearest examples of AI deployment. Every time a card payment is approved online, there may be a fraud system evaluating it in the background. This is especially important for card-not-present transactions, which means online or phone payments where the card is not physically tapped, inserted, or swiped.

Visa, Mastercard, PayPal, Stripe, Square, and other payment companies all describe AI or machine learning as part of their fraud and risk products. Visa says its Visa Protect suite uses AI-powered fraud prevention, and its public materials describe tools such as Visa Advanced Authorization and Visa Deep Authorization that score transactions using hundreds of signals. Because those details come from Visa's own product pages, they should be treated as vendor-published descriptions rather than independent proof of performance. (Visa) Visa also disclosed in an SEC filing that it acquired Featurespace, a company focused on real-time AI payments protection, for $946 million in December 2024. (SEC)

Mastercard's 2025 Form 10-K describes security solutions for prevention, identification, detection, and remediation. It also says its detection systems scan billions of data points across millions of transactions per day and use AI, data analytics, and cyber risk assessment to reduce fraud and false declines. (SEC) Mastercard separately described Decision Intelligence Pro as using generative AI techniques to help assess whether a transaction is valid. (SEC)

For smaller businesses, these systems are often embedded inside tools they already use. Stripe Radar, for example, is a vendor-published fraud tool that says it scores payments using signals from Stripe's network. (Stripe) Square describes risk evaluations based on proprietary machine learning models informed by signals from its payments ecosystem. (Square) Shopify's fraud analysis tools help merchants review orders that may be fraudulent, even though the product page does not necessarily mean each part is advanced AI. (Shopify Help Center)

The important practical point is that many small businesses use AI fraud detection without building AI themselves. A coffee shop, online boutique, repair business, or local nonprofit may rely on fraud scoring built into its payment processor, e-commerce platform, email provider, or bank.

4.4 How AI is used for account security and identity checks

Fraud is not only about stolen cards. Criminals also try to take over accounts, create fake accounts, and impersonate real people. Account takeover means someone gains access to an account that does not belong to them. Synthetic identity fraud means criminals combine real and fake personal information to create a new identity that looks legitimate.

AI tools are used to detect suspicious account behavior. A system might notice that a customer who usually logs in from New Jersey on an iPhone is suddenly logging in from a new device in another country, changing the password, adding a new bank account, and attempting a large transfer. Each action might be allowed by itself, but together they create a risky pattern.

Businesses also use AI in identity verification. Computer vision, a type of AI that analyzes images or video, can compare a selfie to a government ID, detect whether an ID image appears altered, or check whether a person appears to be physically present rather than using a photo. Natural language processing, or NLP, can analyze text such as emails, chat messages, or documents. NLP is AI for processing human language.

This area has become more complicated because criminals are also using AI. FinCEN, the Financial Crimes Enforcement Network, warned financial institutions about fraud schemes involving deepfake media created by generative AI, including fake identity documents used to bypass verification checks. (FinCEN.gov) A deepfake is synthetic audio, video, or imagery that makes it look or sound as if a real person said or did something they did not.

That creates an arms race. Businesses use AI to detect fake IDs and unusual behavior, while criminals use AI to create more convincing fake documents, voices, videos, and messages.

4.5 How AI is used in anti-money laundering and financial crime monitoring

Anti-money laundering, often shortened to AML, refers to systems that try to prevent criminals from disguising illegal money as legitimate money. Know your customer, or KYC, refers to identity checks businesses perform to understand who their customers are. A suspicious activity report, or SAR, is a report that financial institutions may file with regulators when activity appears suspicious.

AI can help with AML because illegal financial activity may involve many small steps across many accounts. A person may move money through several accounts, companies, wallets, or payment methods to hide its source. A single transaction may look normal, but the overall pattern may be suspicious.

The GAO describes financial institutions using AI to assess customer risk, detect illicit activity, screen against sanctions and other lists, analyze transaction data, and review unstructured data such as email, text, or audio for possible money laundering, terrorist financing, bribery, tax evasion, insider trading, and market manipulation. (GAO Files) Unstructured data means information that does not fit neatly into rows and columns, such as a call transcript or email thread.

A common AI method here is graph analytics. A graph is a network of connected items. In fraud detection, the items might be people, accounts, phone numbers, addresses, devices, credit cards, and companies. The connections might show who used the same device, who shared an address, or which accounts transferred money to each other. Graph analytics helps investigators see hidden relationships.

For example, one new account may not look suspicious. But if it shares a device, phone number, shipping address, and bank account with ten previously banned accounts, the network tells a different story.

4.6 How AI is used in cybersecurity

Security teams also use AI to detect attacks on systems and networks. A security information and event management system, or SIEM, collects logs from many systems and helps analysts investigate alerts. A log is a record of what happened, such as a login attempt, file change, or network connection. Endpoint detection and response, or EDR, monitors endpoints such as laptops, servers, and workstations. Extended detection and response, or XDR, combines signals from endpoints, email, cloud systems, identity systems, and networks.

AI helps because cybersecurity produces enormous amounts of data. A large organization may see millions or billions of events per day. Most are harmless. The hard part is finding the few that matter.

In practice, AI security tools may detect suspicious login patterns, identify malware-like behavior, group related alerts into one incident, summarize what happened, recommend next steps, or help analysts write search queries. CrowdStrike's 2026 Form 10-K describes its Falcon platform as AI-native and says it uses security and enterprise data to support real-time detection, investigation, and response. (SEC) Microsoft's Security Copilot documentation describes a generative AI-powered security solution for incident response, threat hunting, intelligence gathering, posture management, and other security tasks. (Microsoft Learn)

Generative AI is AI that creates new content, such as text, code, images, or summaries. A large language model, or LLM, is a generative AI model trained on large amounts of text and code. In security work, an LLM might summarize a phishing investigation, explain a suspicious script, draft an incident report, or translate a natural-language question into a query that searches security logs.

This can be useful, especially for overloaded teams. But it is not the same as fully automated security. Microsoft's own documentation says Security Copilot returns responses for users to review and assess. (Microsoft Learn) That review step matters because AI systems can make mistakes.

4.7 Attackers use AI too

The same AI tools that help defenders can also help attackers. Criminals can use generative AI to write more convincing phishing emails, translate scams into many languages, generate fake customer support messages, create malicious code, or imitate voices.

Verizon's 2026 Data Breach Investigations Report says 31% of breaches started with software vulnerabilities, 48% involved ransomware, and 15 attack techniques were being bolstered by generative AI. (Verizon) Google Threat Intelligence Group reported that threat actors were increasingly integrating AI

into reconnaissance, social engineering, and malware development in late 2025. (Google Cloud) Reconnaissance means gathering information about a target before attacking it. Social engineering means tricking people, rather than only attacking machines.

This is important for students to understand: AI is not automatically "good" or "bad." It is a capability. Defenders use it to find suspicious patterns faster. Attackers use it to produce more convincing deception faster. Businesses therefore need both AI tools and basic security discipline: strong authentication, employee training, software patching, least-privilege access, backups, monitoring, and incident response plans.

4.8 The most common AI techniques

Fraud and security systems use several types of AI. Not every system uses all of them.

Supervised machine learning learns from labeled examples. Labeled means the training data includes the correct answer, such as "fraud" or "legitimate." If a company has years of payment history, it can train a model on past transactions that were later confirmed as fraud or not fraud.

Anomaly detection looks for activity that is unusual compared with normal behavior. An anomaly is something that stands out. A login at 3 a.m. from a new country may be anomalous for one person but normal for another.

Graph analytics studies relationships. This is useful when fraudsters create many accounts that appear separate but share hidden connections.

Computer vision analyzes images or video. It can support ID checks, document review, and liveness detection.

Natural language processing analyzes text or speech. It can help identify phishing emails, suspicious chat messages, fake reviews, or risky customer support conversations.

Generative AI and LLMs help with investigation and explanation. They can summarize alerts, draft reports, explain code, or help analysts search logs. They are usually strongest as assistants, not final decision-makers.

4.9 What this looks like in different-sized businesses

A small business usually does not hire a team of data scientists to build a fraud model. Instead, it uses AI indirectly through services such as payment processors, banks, e-commerce platforms, email security tools, accounting platforms, and cloud security products. The owner may see a simple dashboard: "high-risk order," "possible phishing email," or "unusual login."

A medium-sized business may have more control. It might configure fraud rules, choose risk thresholds, use a SIEM, require multi-factor authentication, and review weekly fraud reports. Multi-factor authentication, or MFA, means users must provide more than one proof of identity, such as a password plus a mobile app approval.

A large business may combine vendor tools with custom models. A bank, marketplace, insurance company, airline, or telecom provider may have a fraud operations team, cybersecurity analysts, data engineers, model-risk staff, and legal or compliance reviewers. These organizations may build models using their own data because they see patterns that no outside vendor can fully see.

In all three cases, the pattern is similar: AI helps prioritize attention. It does not remove the need for policy, judgment, or accountability.

4.10 Risks and limitations

AI fraud and security tools make two basic kinds of mistakes. A false positive happens when the system flags something innocent as suspicious. A real customer's card gets declined, an account is frozen, or an employee is locked out. A false negative happens when the system misses a real threat. A fraudulent order ships, a stolen login succeeds, or malware spreads.

Both are costly. Too many false positives frustrate customers and waste staff time. Too many false negatives allow losses and breaches. The goal is not "perfect AI." The goal is a practical balance between catching risk and allowing normal business to continue.

There are also fairness and privacy concerns. A fraud system may use data that correlates with geography, income, device type, language, or shopping behavior. If not tested carefully, it may treat some groups unfairly. PayPal's SEC risk disclosures, for example, warn that AI and machine-learning algorithms may be flawed or based on biased or insufficient datasets. (SEC)

Another limitation is model drift. Drift happens when the world changes and a model becomes less accurate. Fraud tactics change, customer habits change, products change, and new attack methods appear. A model trained on last year's fraud may not catch this year's fraud unless it is monitored and updated.

Finally, there is an automation risk. A business may be tempted to let the system decide everything. That can be dangerous in high-stakes cases, such as account freezes, denied access, or law-enforcement referrals. The GAO notes that financial regulators using AI generally described AI outputs as informing staff decisions rather than serving as the sole decision-maker. (GAO) NIST's AI Risk Management Framework and its generative AI profile emphasize identifying and managing AI risks rather than assuming AI systems are automatically trustworthy. (nist.gov)

4.11 Hands-on lab: Build a simple fraud risk score in a spreadsheet

This lab shows the basic logic behind many fraud systems. It is not production AI. It is a transparent scoring model, which means you can see exactly how the score is calculated. Real machine-learning systems often learn the weights from historical data, but the idea is similar: combine signals, produce a score, and decide what to do.

Create a spreadsheet with these columns:

txn_id	amount_usd	billing_shipping_mismatch	new_device	failed_login_count	account_age_days	rush_shipping	known_outcome
T001	42	Yes	No	0	365	No	Legit
T002	980	No	Yes	4	2	Yes	Fraud
T003	120	Yes	Yes	0	45	No	Legit
T004	760	Yes	Yes	0	700	No	Legit
T005	215	No	No	1	20	Yes	Fraud
T006	35	Yes	No	0	10	No	Legit

txn_id	amount_usd	billing_shipping_mismatch	new_device	failed_logins_1h	account_age_days	rush_shipping	known_outcome
T007	640	No	Yes	2	3	No	Fraud
T008	88	Yes	Yes	5	1	Yes	Fraud
T009	510	Yes	No	0	180	Yes	Legit
T010	130	No	Yes	0	5	No	Fraud
T011	275	Yes	Yes	4	400	No	Legit
T012	1500	No	No	0	900	No	Fraud

Add a column called risk_score. In row 2, enter this formula and copy it down:

=IF(B2>500,25,0)+IF(C2="No",20,0)+IF(D2="Yes",20,0)+IF(E2>=3,20,0)+IF(F2<7,10,0)+IF(G2="Yes",5,0)

In plain English, the formula says:

Add 25 points if the amount is above $500. Add 20 if the billing and shipping addresses do not match. Add 20 if the customer is using a new device. Add 20 if there were at least three failed logins in the last hour. Add 10 if the account is less than seven days old. Add 5 if rush shipping was chosen.

Now add a column called decision. Use this formula:

=IF(I2>=50,"Review/Block","Approve")

This means any transaction with a risk score of 50 or higher is sent to review or blocked. Anything below 50 is approved.

Finally, compare the decision to known_outcome. A true positive is fraud that was correctly reviewed or blocked. A true negative is a legitimate transaction that was correctly approved. A false positive is a legitimate transaction that was blocked or reviewed. A false negative is fraud that was approved.

Try changing the threshold from 50 to 40, then to 60. Watch what happens. A lower threshold usually catches more fraud but creates more false positives. A higher threshold usually reduces customer friction but lets more fraud through.

That tradeoff is one of the central problems in real fraud detection.

4.12 What students should remember

AI tools are being used in fraud detection and security right now, not just in future predictions. They are used by payment networks, banks, fintech companies, merchants, cybersecurity vendors, cloud platforms, and regulators. Their main job is to detect patterns, assign risk scores, prioritize alerts, and help people respond faster.

But AI is not magic. It can be fooled. It can make unfair or costly mistakes. It can become outdated. It can also be used by attackers. The best real-world systems combine AI with rules, human review, strong authentication, good data governance, monitoring, and clear accountability.

The most realistic way to describe AI in fraud and security is this: AI is a powerful sorting and pattern-recognition tool in a much larger defense system. It helps businesses see danger sooner, but people still have to decide what risks are acceptable, how customers should be treated, and what happens when the system is wrong.

5. AI in Supply Chain Management and Logistics

Course Outcome

VCCS-5. Supply Chain Optimization Using AI Tools: Explore AI's role in supply chain optimization, including demand forecasting, inventory management, and logistics efficiency.

Artificial intelligence tools are already being used in supply chain management and logistics. The more useful answer is that they are usually **not** replacing the whole supply chain. Instead, AI tools are being inserted into specific work steps: forecasting demand, routing deliveries, assigning warehouse tasks, reading freight emails, classifying customs documents, predicting shipment delays, and helping managers decide what to do next.

That distinction matters. A supply chain is the network of people, companies, systems, facilities, and transportation links that move goods from raw materials to customers. Logistics is the part of the supply chain focused on moving, storing, and delivering goods. In real businesses, supply chains are messy. A shipment may depend on weather, labor availability, port congestion, tariffs, fuel prices, warehouse capacity, supplier reliability, customer demand, and software systems that do not always share clean data. AI helps with parts of that mess, but it does not magically make uncertainty disappear.

A good 2026 reality check is this: AI in supply chain management is mostly **embedded**. That means the AI is built into software people already use, such as enterprise resource planning systems, warehouse management systems, transportation management systems, carrier platforms, and customer service tools. Reuters Practical Law described current supply-chain AI use cases as including demand forecasting, procurement and supplier risk, inventory and warehousing, transportation and logistics, customs compliance, fraud prevention, sustainability, and legal support. (Reuters)

Large companies are also telling investors that AI is part of their operating systems, not just an experiment. Walmart's 2026 Form 10-K says the company continues to invest in supply-chain automation and fulfillment and delivery capabilities, and it reports hundreds of distribution facilities across its U.S. and international operations. (Walmart Inc.) UPS tells investors it uses advanced and emerging technologies, including artificial intelligence, for shipment creation, tracking, data management, data-analysis automation, automated agents, personalization, and pricing. (SEC) DHL's 2025 annual report says robotics and

artificial intelligence are increasingly contributing to quality and efficiency in customer service, customs clearance, fulfillment, and service logistics.

The key lesson is that AI is not one thing. In supply chains, it is a family of tools used for prediction, automation, pattern recognition, language processing, and optimization.

5.1 The basic AI terms students need

Artificial intelligence, or **AI**, means computer systems that perform tasks that normally require human judgment, such as recognizing patterns, predicting outcomes, generating text, or recommending actions. In business, AI usually works by finding patterns in data and using those patterns to support a decision.

Machine learning is a type of AI where software learns from examples instead of being programmed with every rule by hand. For example, a delivery company may train a model on past delivery data so it can estimate whether today's package will arrive late.

Generative AI is AI that creates new content, such as text, summaries, emails, code, or images. A large language model, or **LLM**, is a generative AI model trained to work with language. In logistics, an LLM might read a shipper's email, extract the pickup location and delivery date, and draft a response.

Computer vision is AI that analyzes images or video. In a warehouse, computer vision can help identify damaged packages, read labels, check whether a pallet is loaded correctly, or guide a robot arm.

Optimization is the process of finding the best choice under constraints. A route-planning system, for example, may try to minimize miles driven while also respecting delivery windows, truck capacity, driver hours, traffic, and customer preferences. Optimization is not always AI by itself, but modern logistics tools often combine optimization with machine learning predictions.

Agentic AI refers to software agents that can take a sequence of steps toward a goal. In business settings, an AI agent might read an email, check a shipment record, update a transportation system, and draft a reply. In serious deployments, these agents are usually limited by permissions, audit logs, and human review for higher-risk decisions.

5.2 Demand forecasting: predicting what customers will need

One of the most common supply-chain uses of AI is **demand forecasting**, which means predicting future customer demand. A forecast might answer questions like: How many cases of bottled water will a store need next week? How many replacement parts should a manufacturer keep in stock? Which products should a retailer move closer to customers before a holiday weekend?

Traditional forecasting used historical sales data and simple formulas. AI forecasting can combine many more signals: past sales, seasonality, promotions, local events, weather, e-commerce searches, price

changes, economic conditions, and product substitutions. The goal is not perfect prediction. The goal is better decisions about purchasing, production, inventory, labor, and transportation.

This matters because two costly problems sit on opposite sides of the same decision. A **stockout** happens when a company runs out of a product customers want. An **overstock** happens when the company buys or produces too much and has to store, discount, or discard it. Better forecasts can reduce both, although no forecast can eliminate uncertainty.

Walmart has publicly described its newer supply-chain systems as using real-time AI and automation to predict demand, reroute inventory, reduce waste, and simplify work across markets including Costa Rica, Mexico, and Canada. Because this is company-published information, it should be read as Walmart's own description of its system, not as an independent audit of results. (Walmart News & Leadership) Amazon also describes demand forecasting as one of the practical AI applications used in its logistics network. (Amazon News)

For students, the important concept is that forecasting is not just "guessing sales." It changes real operations. A forecast may trigger a purchase order, move inventory from one warehouse to another, schedule workers, reserve truck capacity, or warn a manager that a product is at risk of running out.

5.3 Transportation and last-mile delivery: planning routes and predicting exceptions

Transportation is another major AI use area. The **last mile** is the final movement of a product from a distribution point to the customer's door or business location. It is expensive and difficult because it involves traffic, parking, delivery windows, missing apartment numbers, theft risk, bad weather, and customers changing plans.

AI tools are used to estimate delivery times, design routes, predict failed deliveries, identify risky addresses, and recommend changes when conditions shift. Business Insider reported in 2025 that UPS's ORION routing system, originally launched in 2013, has evolved into a machine-learning tool that can shorten routes and reroute drivers based on changing conditions. (Business Insider) UPS also reports in its own filings that it uses AI and digital tools for shipment creation, tracking, data management, automated agents, and pricing. (SEC)

This does not mean a delivery driver simply follows whatever an AI says. The route plan is a decision aid. The driver still deals with real-world constraints: a locked gate, a blocked loading dock, an unsafe stop, a customer who is not home, or a road closure that the system has not recognized yet.

Amazon has also described AI tools for delivery accuracy. Its "Wellspring" system is presented by Amazon as using signals from maps, addresses, delivery history, and other data to improve delivery-location accuracy. (Amazon News) Again, the source is Amazon's own publication, so the safest interpretation is that these tools show where the company is deploying AI, not that every performance claim has been independently verified.

The business value is clear: even a small improvement in miles driven, failed delivery rates, or package placement accuracy can matter at massive scale. But the risk is also clear. A routing model can be wrong, and a wrong model can create safety issues, unfair workloads, or poor customer service if no one checks its effects.

5.4 Warehouses and fulfillment: robots, cameras, and task assignment

Warehouses are among the most visible places where AI touches physical work. A warehouse management system, or **WMS**, is software that helps control receiving, storage, picking, packing, and shipping. AI can be added to a WMS to decide which worker or robot should handle a task, which items should be stored near each other, or which orders should be picked together.

Modern fulfillment centers use several kinds of automation. Some robots move shelves or containers. Others sort packages, scan labels, or assist workers. Computer vision can inspect items, read barcodes, or detect exceptions. AI can also predict congestion inside a facility and adjust task assignments.

Amazon is one of the clearest examples because it operates at enormous scale and publishes technical descriptions of its robotics systems. In 2025, Amazon said it had deployed more than one million robots across its operations and introduced DeepFleet, a generative AI foundation model designed to coordinate robot movement and improve robot fleet travel time. (Amazon News) Amazon Science described DeepFleet as being trained on millions of hours of fulfillment and sort-center data and as predicting interactions among mobile robots so the system can assign tasks and route around congestion. (amazon.science)

The phrase **foundation model** means a large AI model trained on broad data that can be adapted to many related tasks. In this case, the model is not writing poetry or chatting with customers. It is helping coordinate physical movement inside logistics facilities.

DHL's annual report gives a broader logistics-industry view. DHL says robotics and AI are contributing to improved quality and efficiency in fulfillment, customer service, customs clearance, and service logistics. That phrasing is useful because it is measured and operational. It does not claim that warehouses are "fully autonomous." It says AI and robotics are becoming levers for quality and efficiency.

For students, the main point is that warehouse AI is usually connected to work queues. It asks: What task should happen next? Who or what should do it? Where is the item? What is the fastest safe path? Is something unusual happening?

5.5 Freight brokerage and logistics paperwork: where generative AI is very practical

Some of the most realistic AI deployments are not flashy robots. They are tools that read emails, documents, forms, and status messages.

A **freight broker** helps match shippers that need to move goods with carriers that can move them. A **third-party logistics provider**, or **3PL**, manages logistics services for another company. These businesses handle huge volumes of repetitive communication: quote requests, pickup appointments, load tenders, tracking updates, proof-of-delivery documents, invoices, and exception notices.

This is where generative AI can be useful. Many logistics processes still run through email and PDFs. A shipper may email, "Can you move 18 pallets from Atlanta to Dallas next Tuesday?" A human employee or AI tool must extract the origin, destination, date, weight, equipment type, and delivery requirements, then check prices and capacity.

C.H. Robinson, a large logistics provider, has described using generative AI agents to automate parts of the shipment lifecycle. In 2025, the company said its AI had performed more than three million shipping tasks, including over one million price quotes and over one million orders processed. It also said its agents

can quote prices, process orders, acquire trucking capacity, set appointments, check loads in transit, and provide tracking updates. These are company-published numbers, so they should be treated as disclosed corporate claims rather than independent evaluation. (C.H. Robinson)

In a 2024 company release, C.H. Robinson said generative AI was reading emails and automating tasks such as price quotes, load acceptance, scheduling pickup and delivery appointments, and checking loads. The company said more than 10,000 email transactions per day no longer required manual handling. (C.H. Robinson)

This is a powerful example because it shows what AI agents often do in real business: they sit between messy human communication and structured business systems. They convert unstructured information, such as an email, into structured fields, such as pickup date, delivery address, pallet count, and carrier rate. Then they either update a system or prepare an action for a human to approve.

DHL has also discussed autonomous AI agents for logistics communication. In 2025, DHL described working with HappyRobot on AI agents that can interact through phone, email, and messaging and integrate with DHL internal systems. (DHL Group) That does not mean every call center or dispatcher disappears. It means companies are testing or deploying AI on routine interactions where speed and consistency matter.

5.6 Customs, tariffs, and trade compliance

International supply chains create a different kind of challenge: rules. Goods crossing borders may need tariff classification, export-control screening, customs declarations, country-of-origin documentation, product descriptions, and restricted-party checks. A mistake can delay shipments or create legal penalties.

A **tariff** is a tax or duty on imported goods. A **customs classification code** is a standardized code used to identify a product for customs and tariff purposes. In the United States and many other countries, product classification can be extremely detailed. Similar-looking products may have different duty rates depending on material, use, and origin.

AI can help by reading product descriptions, suggesting classification codes, checking documents for missing information, and flagging risky shipments. It does not remove legal responsibility. A company still needs compliance professionals, especially for high-value, regulated, or ambiguous goods.

Maersk, one of the world's largest shipping and logistics companies, rolled out an AI-backed customs platform called Trade & Tariff Studio in 2025. FreightWaves reported that the system was designed to help with complex customs processes, product codes and subcodes, risk screening, and tariff optimization, with support from customs experts and data partners. (FreightWaves) Oracle's supply-chain software documentation also lists AI-assisted features for global trade management, including harmonized-system code classification and trade compliance tasks. (Oracle Docs)

This is a good example of AI as "first-pass assistant." The model may suggest a code or flag an issue, but a human expert may still need to approve the classification, especially when the financial or legal stakes are high.

5.7 Supplier risk, disruption management, and control towers

A supply chain manager does not only ask, "Where is my shipment?" They also ask, "What might go wrong next?"

A **supplier risk** tool helps evaluate whether a supplier may fail to deliver because of financial trouble, quality problems, labor disruption, geopolitical risk, severe weather, or regulatory changes. A **control tower** is a dashboard that gives managers visibility across orders, inventory, shipments, suppliers, and exceptions. The name sounds dramatic, but in practice it usually means software that combines data from multiple systems and highlights issues.

AI can scan shipment data, supplier performance, weather alerts, port delays, news, and internal orders to detect risks. It may recommend expediting a shipment, switching suppliers, moving inventory, or warning a customer earlier. Reuters Practical Law described AI supply-chain risk use cases as including supplier monitoring, contract review, sanctions screening, and risk assessment. (Reuters)

This area also shows why "more AI" is not automatically better. If a model is trained mostly on normal conditions, it may struggle with rare disruptions: a major war, sudden tariff change, pandemic, cyberattack, bridge collapse, port strike, or extreme weather event. Supply chains need human judgment because many disruptions are unusual, political, or relationship-based. A supplier might be late not because it is unreliable, but because a customer changed specifications or because a port inspection delayed containers.

The best control-tower tools therefore combine prediction with escalation. The system surfaces the problem; a planner decides what to do.

5.8 What the technology stack looks like

A supply-chain AI system usually has four layers.

The first layer is the **data layer**. It includes orders, invoices, purchase orders, inventory records, barcode scans, GPS pings, RFID tags, carrier messages, customer-service tickets, supplier data, weather data, and customs documents. RFID, or **radio-frequency identification**, uses small tags and radio signals to identify items without scanning each barcode directly.

The second layer is the **model layer**. This is where machine learning models, forecasting models, optimization engines, computer-vision tools, and large language models do their work. One model might predict demand. Another might estimate arrival time. Another might classify a customs document.

The third layer is the **workflow layer**. This is where employees actually interact with the system: an alert in a transportation dashboard, a recommended reorder quantity in an inventory screen, a suggested supplier email, a robot task assignment, or a customer-service response.

The fourth layer is the **governance layer**. Governance means the rules, controls, and accountability around a system. Who can use the AI? What data can it access? Which decisions require human approval? Are outputs logged? Can the company explain why a shipment was rejected, a supplier was flagged, or a customer was charged a certain rate?

This structure helps explain why AI adoption is uneven. A large retailer or carrier may have enough data, engineers, and process discipline to build or customize AI. A small business may use AI only through tools already built into software from a vendor, carrier, marketplace, or 3PL.

That vendor embedding is accelerating. Oracle announced AI agents for Oracle Fusion Cloud Applications in 2025, including agents for supply-chain planning, fulfillment, and process automation. (Oracle) SAP announced Joule agents for supply-chain tasks such as production-order checks, change management, and supplier onboarding, with some availability planned in 2026. (SAP News Center) These announcements do not prove customer results by themselves, but they show that major enterprise software vendors are making AI part of standard business applications rather than a separate experimental tool.

5.9 What is real, what is exaggerated

The real part is that AI is already changing specific supply-chain jobs. Planners get better forecast suggestions. Dispatchers get route recommendations. Warehouse workers interact with robot-assisted systems. Freight coordinators spend less time copying information from emails into transportation systems. Customs teams receive suggested classifications and risk flags. Customer-service agents get summaries and recommended replies.

The exaggerated part is the idea of a fully autonomous supply chain that runs itself from supplier negotiation to final delivery. That is not the normal reality in 2026. Most companies still have fragmented data, legacy systems, supplier exceptions, contract complexity, and human relationships. Many still depend on spreadsheets, email, phone calls, and manual workarounds.

There is another reason full autonomy is limited: accountability. If an AI system chooses the wrong supplier, misclassifies a product, violates a delivery rule, or creates an unsafe route, the company cannot blame "the algorithm" and walk away. A business still has legal, financial, safety, and customer obligations.

The best way to understand AI in logistics is not "human versus machine." It is **human plus machine plus process**. A weak process with AI added to it may simply fail faster. A strong process with good data and clear controls may become faster, more reliable, and easier to manage.

5.10 Risks and governance

AI in supply chains creates several serious risks.

The first is **data quality**. If inventory records are wrong, supplier names are inconsistent, delivery scans are missing, or product descriptions are vague, the AI may produce confident but incorrect recommendations. In logistics, bad data can cause real-world costs: a truck goes to the wrong dock, a warehouse runs out of space, or a customer is promised inventory that does not exist.

The second risk is **hallucination**, a term used when generative AI produces information that sounds plausible but is false. In a supply chain, a hallucinated delivery promise, tariff code, or supplier instruction could be expensive. This is why generative AI outputs need verification when they affect money, law, safety, or customer commitments.

The third risk is cybersecurity. UPS explicitly warns investors that increased AI use can increase cybersecurity risks because AI systems may process sensitive data or be targeted by malicious actors. (SEC) A logistics system may contain customer addresses, shipment contents, pricing, supplier contracts, customs documents, and operational schedules. That information is valuable.

The fourth risk is unfair or unsafe optimization. A route model might create unrealistic driver workloads. A warehouse algorithm might assign tasks in a way that increases injury risk. A supplier-risk model might penalize small suppliers because they have less digital data, not because they perform worse.

The fifth risk is overdependence. When companies optimize tightly for cost and speed, they may remove backup capacity. That can make the network fragile when disruption occurs. AI should support resilience, not only efficiency.

Governance frameworks are emerging to help organizations manage these risks. The National Institute of Standards and Technology, or NIST, created the AI Risk Management Framework to help organizations manage AI risks, and NIST's 2024 generative AI profile focuses on risks and controls specific to generative AI. (nist.gov) The European Union's AI Act entered into force in 2024 and becomes fully applicable in stages, with broad applicability by August 2026, which matters for companies operating in or selling into the European market. (Digital Strategy EU)

For a student entering business or IT, the practical question is not just "Can we use AI?" It is "What decision is this AI changing, what could go wrong, and who is responsible for checking it?"

5.11 Hands-on lab: build a simple AI-supported inventory decision

This lab uses a small example so you can see the logic behind a real supply-chain decision. You can do it in Excel, Google Sheets, or any spreadsheet. You may also use a chatbot as an assistant, but the spreadsheet math is the source of truth.

5.11.1 Scenario

You manage inventory for a small campus store that sells a popular water bottle. The supplier takes two weeks to deliver after you place an order. You want to decide whether to reorder.

Here is your weekly sales history:

Week	Units sold
1	120
2	132
3	128

Week	Units sold
4	145
5	160
6	170
7	155
8	180

5.11.2 Step 1: Create a simple forecast

A **moving average** forecast uses the average of recent periods to predict the next period. It is not advanced AI, but it teaches the same basic idea: use past data to estimate future demand.

Use a three-week moving average.

Week 4 forecast = average of Weeks 1, 2, and 3 = (120 + 132 + 128) / 3 = 126.7 units

Continue the same method:

Forecast for week	Calculation	Forecast
4	Average of Weeks 1–3	126.7
5	Average of Weeks 2–4	135.0
6	Average of Weeks 3–5	144.3
7	Average of Weeks 4–6	158.3
8	Average of Weeks 5–7	161.7
9	Average of Weeks 6–8	168.3

Your Week 9 forecast is about **168 units**.

5.11.3 Step 2: Add lead time and safety stock

Lead time is how long it takes to receive inventory after ordering. Here, lead time is two weeks.

Safety stock is extra inventory kept as a cushion against uncertainty. Here, use 60 units.

A simple reorder point is:

Reorder point = expected demand during lead time + safety stock

Expected demand during lead time = 168.3 units per week × 2 weeks = 336.6 units

Reorder point = 336.6 + 60 = 396.6, rounded to **397 units**

Plain-English translation: if your inventory falls below 397 units, you may not have enough to cover expected sales during the two-week wait, plus a small cushion.

5.11.4 Step 3: Make the decision

Suppose you currently have 310 units on hand.

Because 310 is below the reorder point of 397, you should reorder.

Now choose a simple target stock level. One practical target is:

Target stock = reorder point + one more week of forecast demand = 397 + 168 = 565 units

Recommended order quantity = target stock − inventory on hand = 565 − 310 = **255 units**

So your recommendation is: order about **255 water bottles**.

5.11.5 Step 4: Use generative AI carefully

Now use a chatbot as an assistant, not as the authority. Paste this prompt:

You are a supply-chain analyst for a small campus store. Weekly sales for a water bottle were 120, 132, 128, 145, 160, 170, 155, and 180 units. I used a three-week moving average and got a Week 9 forecast of 168.3 units. Supplier lead time is two weeks, safety stock is 60 units, and current inventory is 310 units. Check whether I should reorder, explain the reasoning in plain English, list three uncertainties, and draft

a short supplier email asking about availability and lead time. Do not invent discounts, delivery promises, or supplier names.

Then verify the AI's answer. Did it keep the arithmetic correct? Did it avoid inventing facts? Did it explain uncertainty? Did it clearly separate recommendation from assumption?

This is what real AI-supported supply-chain work often looks like. The AI helps summarize, explain, draft, and check. The human still owns the decision.

5.12 What students should take away

AI tools are being used in supply chain management and logistics today, but their value comes from targeted deployment. They forecast demand, route deliveries, coordinate robots, read freight messages, classify documents, detect risks, and help workers respond faster. They are most useful when connected to clean data, clear processes, and responsible human oversight.

The most important career skill is not memorizing one AI product. Products will change. The durable skill is understanding the workflow. Ask: What decision is being made? What data supports it? What system will the AI update? What happens if the AI is wrong? Who reviews the output? How will the company measure whether the tool actually improved cost, speed, service, safety, or resilience?

For supply chains, AI is not a magic brain floating above the business. It is a set of tools inside the business. The companies that benefit are usually the ones that combine AI with operational knowledge, disciplined data, and thoughtful controls.

6. AI Tools in Human Resources

Course Outcome

VCCS-6. AI-Driven Human Resource Management: Assess how AI-driven tools streamline human resource management, including recruitment automation, resume screening, and employee engagement analysis.

Artificial intelligence tools are already being used in **Human Resources**, usually called **HR**. HR is the business function responsible for hiring, onboarding, payroll coordination, benefits, employee records, training, performance reviews, promotions, workforce planning, and sometimes workplace investigations or employee relations. In 2026, AI in HR is not mainly a science-fiction "robot boss." It is more often a set of features built into software that companies already use: applicant tracking systems, payroll platforms, human capital management systems, learning platforms, employee help desks, and workforce scheduling tools.

A useful way to summarize the current reality is this: **AI is being used most heavily where HR work is repetitive, document-heavy, data-heavy, or communication-heavy.** That includes writing job postings, screening résumés, matching candidates to jobs, scheduling interviews, answering employee questions, detecting payroll anomalies, recommending training, summarizing performance feedback, and forecasting staffing needs. It is used much less consistently for final decisions about hiring, firing, promotion, discipline, or pay, partly because those decisions carry legal, ethical, and reputational risk.

A 2026 SHRM survey of HR professionals found that 39% said AI had already been adopted inside their HR function, while another 7% planned to launch HR AI that year. The most common HR areas were recruiting, HR technology, learning and development, and employee experience. SHRM also found that many organizations still do not formally measure whether their AI investments are working, which is important: AI adoption is real, but its value is not always carefully proven. (SHRM)

6.1 Basic Concepts: What "AI in HR" Means

Artificial intelligence, or **AI**, means computer systems that perform tasks associated with human judgment, language, pattern recognition, prediction, or decision support. In HR, AI does not usually "understand people" the way a human manager might. It looks for patterns in data and produces outputs such as recommendations, rankings, summaries, predictions, or generated text.

Machine learning is a type of AI in which software learns patterns from data. For example, a machine learning model might look at past employee turnover data and identify patterns associated with people leaving the company. That does not mean the model knows why someone will quit. It means the system has found statistical relationships in the data.

Generative AI is AI that creates new content, such as text, images, code, or summaries. In HR, generative AI might draft a job description, summarize an employee handbook, write a first draft of a performance review, or create a training quiz.

A **large language model**, or **LLM**, is a generative AI system trained on huge amounts of text. It predicts and generates language. ChatGPT is one example of an LLM-based tool. HR departments may use LLMs directly, but more often they use them through business software platforms that embed generative AI into everyday workflows.

A **human capital management system**, or **HCM system**, is software that helps organizations manage employee information and HR processes. HCM systems may include recruiting, onboarding, payroll, benefits, performance management, learning, internal mobility, and analytics.

The most important point for students is that AI in HR is usually not one single product. It is a layer inside many products. A company might not say, "We bought an AI HR system." It might say, "We use Workday, ADP, Dayforce, Oracle, SAP, UKG, Greenhouse, iCIMS, LinkedIn Recruiter, or another HR platform," and those systems may include AI features.

6.2 Where AI Is Most Common: Recruiting and Hiring

The most visible use of AI in HR is **talent acquisition**, which means finding, attracting, evaluating, and hiring job candidates. This is where HR departments face large volumes of documents and messages. A single job opening can attract hundreds or thousands of applications. AI tools are attractive because they promise to reduce repetitive screening and communication work.

Common recruiting uses include:

- drafting job descriptions and job ads;
- rewriting job postings for clarity or tone;
- parsing résumés into structured fields;
- matching candidate skills to job requirements;
- ranking or sorting applicants;
- answering candidate questions through chatbots;
- scheduling interviews;

- generating interview questions;
- summarizing interview notes;
- helping recruiters search internal or external talent databases.

SHRM's 2024 research found that among organizations using AI in HR, talent acquisition was the most common area. Among organizations using AI for recruiting, common tasks included generating job descriptions, customizing job postings, reviewing or screening résumés, communicating with applicants, and automating candidate searches. (SHRM)

Here is a realistic example. A retailer needs to hire 200 seasonal warehouse workers. The HR team posts jobs in multiple cities. AI helps generate job ads, a chatbot answers applicant questions, the system checks whether applicants meet basic requirements, and scheduling software books interviews or hiring events. A human recruiter may still make the final call, but the workflow is heavily automated.

Large HR platforms are building these features directly into their products. Workday, one of the major enterprise HCM vendors, describes its HCM suite as covering the employee lifecycle from recruiting through retirement. In its 2026 annual filing, Workday described AI-powered talent acquisition products including HiredScore AI for Recruiting, Candidate Experience Agent, and Paradox Conversational Applicant Tracking System. (SEC) ADP, another major payroll and HR vendor, describes AI features in ADP Assist across payroll, time, talent, benefits, recruitment, analytics, reporting, and compliance. Its filings also describe tools that use large skills datasets to support candidate matching and labor market insight. (SEC)

This matters because HR AI is not limited to experimental startups. It is being embedded inside mainstream enterprise software used by large and midsize employers. Smaller employers may also use AI without building anything themselves, because AI features can appear inside payroll, scheduling, job posting, or applicant tracking tools they already subscribe to.

6.2.1 The Hiring Risk: Bias at Scale

Hiring is also the area where AI risk is easiest to understand. If a human recruiter unfairly screens out candidates, that is a serious problem. If software unfairly screens out candidates at scale, the problem can affect thousands of people before anyone notices.

A key legal concept is **adverse impact**, sometimes called **disparate impact**. This means a seemingly neutral practice may disproportionately exclude people in a protected group, such as by race, sex, age, disability, religion, or national origin. The U.S. Equal Employment Opportunity Commission, or **EEOC**, has said that federal employment discrimination law applies when employers use automated systems to make or inform employment selection decisions. The EEOC's guidance discusses the "four-fifths rule," a rule of thumb used to flag selection rates that may need closer review. (EEOC)

The risk is not theoretical. In one EEOC case, iTutorGroup was accused of programming online recruitment software to automatically reject female applicants age 55 or older and male applicants age 60 or older. The company later agreed to pay $365,000 to settle the lawsuit. (EEOC) (EEOC)

The lesson is simple: **automating a bad rule does not make it fair.** AI can make HR faster, but it can also make unfair screening faster.

6.3 Employee Help Desks, Payroll, and Benefits

A second major use of AI in HR is employee support. Many HR departments answer the same questions repeatedly:

"Where do I find my W-2?" "How many vacation days do I have?" "What is the parental leave policy?" "How do I change my direct deposit?" "What happens if I miss open enrollment?"

AI chatbots and generative AI assistants can answer common employee questions by searching company policy documents, benefit materials, payroll information, and HR knowledge bases. This type of AI is often less glamorous than résumé screening, but it can be very useful. HR teams spend a lot of time responding to routine questions, and employees often need answers outside normal office hours.

ADP's 2024 annual filing gives a concrete example. ADP Assist uses generative AI across payroll, time, talent, benefits, recruitment, analytics, reporting, and compliance. ADP describes uses such as validating payroll information, checking payroll anomalies, identifying missing tax registrations, answering compliance questions, simplifying report creation, and helping users access workforce data. (SEC)

This kind of AI usually works best when it is connected to reliable internal documents. A chatbot that answers from the actual employee handbook is more useful than a chatbot that guesses. But there is still risk. A chatbot could give an outdated answer about medical leave, misstate a benefit rule, or reveal

information to the wrong person. For that reason, many organizations treat HR chatbots as assistants, not final authorities.

A responsible HR chatbot should be able to say, "I do not know," cite the policy it used, escalate to a human HR representative, and avoid exposing private employee data.

6.4 Learning, Development, and Skills

Another growing use of AI in HR is **learning and development**, often abbreviated as **L&D**. This means helping employees build skills through training, coaching, certifications, courses, mentoring, and career pathways.

AI can help L&D teams in several ways. It can recommend training based on a worker's role, career goals, or skill gaps. It can generate quizzes, practice scenarios, simulations, and course outlines. It can summarize training materials or convert a long policy document into a shorter lesson. It can also help build a **skills inventory**, which is a structured list of skills employees have or may need.

SHRM's 2026 report found learning and development among the more common HR AI areas, after recruiting and HR technology. It also described AI uses such as content generation, AI-generated quizzes and scenarios, candidate-job matching, and personalized learning recommendations. (SHRM)

The idea of skills is especially important. Many companies want to move from job-title-based HR to **skills-based talent management**. Instead of saying, "Maria is a payroll specialist," the system might say, "Maria has skills in payroll compliance, Excel, Spanish-language customer support, and tax reporting." That can help companies find internal candidates for projects or promotions.

But skills data can be incomplete. Some employees are better at describing their work than others. Some managers document employee skills carefully; others do not. Some skills are visible in software systems, while others are informal or relational. If AI recommends opportunities only to people whose skills are already well documented, it may reinforce existing inequalities.

A good HR team should ask: Who is missing from the data? Whose work is easy to measure? Whose work is invisible to the system?

6.5 Performance Reviews and Manager Support

AI is also being used in **performance management**, the HR process for setting goals, giving feedback, evaluating work, and documenting performance. Generative AI can help managers write clearer feedback, summarize peer comments, draft goals, and turn rough notes into more professional language.

This can be helpful because many managers struggle to write useful performance reviews. Some write vague comments such as "good team player" or "needs improvement." AI can suggest more specific wording, remind managers to connect feedback to goals, or identify missing examples.

SHRM's 2024 research found that among organizations using AI for performance management, many used it to help managers provide more comprehensive and actionable feedback or to support employee goal setting. (SHRM)

But performance reviews are sensitive. A poorly written AI-generated review can damage someone's career. AI may produce confident but unsupported language. It may soften serious issues too much or make minor issues sound severe. It may also reproduce patterns in prior reviews, including biased language. For example, if women in past reviews were more often described as "helpful" while men were described as "strategic," an AI tool trained on those patterns could repeat them.

The safest use is not "let AI rate the employee." The safer use is "let AI help the manager organize evidence, check clarity, and draft language that the manager must verify." The human manager remains responsible for accuracy and fairness.

6.6 Workforce Planning and People Analytics

HR departments also use AI for **people analytics**, which means analyzing workforce data to support business decisions. Examples include turnover, absenteeism, pay, promotion patterns, hiring speed, engagement survey results, training completion, and staffing levels.

AI can help answer questions such as:

How many nurses will this hospital need next quarter? Which departments have unusually high turnover? Which skills are becoming harder to hire for? Are employees leaving after a certain manager, schedule, or commute pattern? How might a hiring freeze affect customer support wait times?

Workday describes AI-supported planning features that assist forecasting using historical and third-party data, including labor statistics. (SEC) ADP also describes analytics tools that help users analyze compensation, turnover, candidate profile relevancy, and talent market insights. (SEC)

People analytics can improve planning, but it must be handled carefully. Employee data is not just business data; it is personal data about people's jobs, pay, schedules, health benefits, locations, performance, and sometimes family situations. A turnover prediction model, for example, might label an employee as a "flight risk." That could help a manager offer support, but it could also unfairly limit the employee's promotion opportunities if the label is misused.

Prediction is not destiny. A model might find that employees with long commutes are more likely to quit. That does not mean any one employee with a long commute is planning to quit. HR teams need to avoid turning probabilities into stereotypes.

6.7 Scheduling, Monitoring, and Algorithmic Management

A more controversial use of AI in HR is **algorithmic management**. Algorithmic management means using software, sometimes including AI, to automate tasks traditionally done by human managers. That can include assigning shifts, measuring productivity, ranking workers, triggering warnings, recommending discipline, or controlling work pace. The OECD describes algorithmic management as software that fully or partly automates managerial tasks, with possible productivity benefits but also risks for workers. (OECD)

This type of AI is common in frontline settings: warehouses, delivery, retail, call centers, restaurants, transportation, and gig-work platforms. AI may forecast customer demand, assign workers to shifts, track call times, monitor route completion, evaluate productivity, or flag unusual behavior.

Some uses are reasonable. A hospital needs enough staff on a shift. A grocery store needs cashiers when customers arrive. A call center needs to know whether customers are waiting too long. AI can help forecast workload and improve scheduling.

But monitoring can become harmful when workers feel constantly watched or when flawed metrics drive discipline. The U.S. Government Accountability Office reviewed research on digital workplace surveillance and noted that employers monitor workers for performance, productivity, and safety, while studies have examined effects on physical health, mental health, and employment opportunities. (GAO) In a related report, GAO noted stakeholder concerns that digital surveillance can create distrust, reduce morale, and discourage workers from exercising workplace rights. (GAO)

The core issue is not simply "AI is bad" or "monitoring is bad." The issue is whether measurement is fair, accurate, transparent, and humane. Measuring keystrokes may be a poor way to evaluate a software developer. Measuring call length may punish a customer support worker who handles complex problems. Measuring warehouse speed without considering safety may encourage injuries.

Good HR technology should help people do better work. It should not reduce people to a dashboard number.

6.8 Employees Often Experience AI Differently Than Executives Do

One reason AI in HR is confusing is that leaders and employees may see different realities. Senior executives may talk about AI strategy, while many workers have not been trained or even told clearly how AI is being used.

Gallup found that 93% of Fortune 500 chief human resources officers said their organization had begun using AI tools or technologies to improve business practices, but only 33% of U.S. employees said their organization had begun integrating AI. Gallup also found that many employees had not received clear guidance or training. (Gallup.com) Pew Research Center similarly found that many U.S. workers were more worried than hopeful about future workplace AI use, and only a minority said at least some of their current work was being done with AI. (Pew Research Center)

This gap matters for HR. HR is often responsible for communication, training, policy, and trust. If employees believe AI is being introduced secretly, or mainly to replace or monitor them, they may resist it. If employees understand what the tool does, what it does not do, what data it uses, and how to challenge errors, adoption is more likely to be legitimate.

6.9 Regulation and Governance Are Becoming Part of HR AI

AI in HR is increasingly regulated, especially when it affects hiring, promotion, pay, discipline, or termination. The legal landscape is changing quickly, but the direction is clear: employers are being pushed to document, test, explain, and govern automated employment tools.

New York City's Local Law 144 is one example. It restricts employer use of automated employment decision tools unless a bias audit has been conducted, certain information is made public, and notices are provided to candidates or employees. (New York City Government) Enforcement is not simple, however. A 2025 New York State Comptroller audit found that the city agency responsible for enforcement had difficulty identifying noncompliance, especially when employers did not disclose AI use or post required audit information. (Office of the New York State Comptroller)

Colorado has also moved into this area. In May 2026, Colorado enacted SB26-189, replacing and revising its earlier AI law. The law defines automated decision-making technology and covers consequential decisions, including employment. Starting January 1, 2027, developers and deployers of covered systems face requirements involving documentation, notices, records, adverse-outcome explanations, correction requests, and meaningful human review. (Colorado General Assembly)

California's Civil Rights Council approved regulations clarifying that employment discrimination protections apply to AI, algorithms, and automated decision systems. (Civil Rights Department) In the European Union, the AI Act treats employment and worker-management AI systems as high-risk, meaning they face stricter requirements around documentation, risk management, data governance, human oversight, accuracy, and cybersecurity. (Digital Strategy EU) (Artificial Intelligence Act EU)

For HR professionals, the takeaway is practical: AI governance is becoming part of HR work. HR teams need to know what systems they use, what decisions those systems affect, what data they process, how bias is tested, how humans review outputs, and how employees or applicants can challenge errors.

6.10 What Responsible HR AI Looks Like

A responsible HR AI program does not start with a flashy tool. It starts with a clear question: **What HR problem are we trying to solve, and what would count as success?**

For example, "We want AI" is not a good goal. "We want to reduce interview scheduling time while maintaining candidate satisfaction and avoiding discriminatory screening" is better. "We want to use AI to identify internal employees who may qualify for cybersecurity training, while allowing people to add missing skills to their profiles" is also better.

Responsible HR AI usually includes several practices.

First, the organization keeps an inventory of AI tools. HR cannot govern tools it does not know about. That inventory should include vendor tools, internally built tools, and general-purpose AI tools used by HR staff.

Second, the organization classifies use cases by risk. Drafting a picnic announcement is low risk. Ranking job applicants is high risk. Recommending termination is very high risk.

Third, the organization tests for accuracy and bias. In hiring, that may include adverse-impact analysis. In payroll, it may include checking whether anomaly detection creates false alarms. In chatbots, it may include testing whether answers match current policy.

Fourth, the organization keeps humans accountable. "Human in the loop" should not mean a person blindly clicks approve. It should mean the human reviewer has enough information, authority, and time to question the system.

Fifth, the organization communicates clearly. Applicants and employees should know when automated tools are being used in important decisions, what data is involved, and how to request review or correction.

Finally, the organization measures outcomes. This is where many companies are still immature. SHRM found that more than half of HR professionals said their organizations did not formally measure the success of AI investments. (SHRM) Without measurement, AI can become expensive decoration: impressive in demos but weak in real impact.

6.11 Hands-On Lab: Run a Mini Bias Audit on an AI Screening Tool

This lab gives you a simplified version of an HR AI governance task. You will test whether an AI screening tool may be creating an adverse-impact warning. This is not a full legal audit. It is a classroom exercise using the four-fifths rule as a practical screening method.

6.11.1 Scenario

A company used an AI-assisted résumé screening tool for an entry-level analyst job. The tool did not make final hiring decisions. It recommended applicants for recruiter review. HR wants to check whether recommendation rates differ sharply across demographic groups.

In a real company, demographic data must be handled carefully and legally. For this lab, we will use anonymous groups: Group A, Group B, and Group C.

6.11.2 Data

Group	Applicants	Recommended for recruiter review
Group A	100	50
Group B	80	28
Group C	40	18

6.11.3 Step 1: Calculate the selection rate

The **selection rate** is the percentage of applicants in a group who were selected or recommended.

Formula:
Selection rate = Recommended applicants ÷ Total applicants
For Group A:
50 ÷ 100 = 0.50, or 50%
For Group B:
28 ÷ 80 = 0.35, or 35%
For Group C:
18 ÷ 40 = 0.45, or 45%

6.11.4 Step 2: Find the highest selection rate

The highest rate is Group A's 50%.

6.11.5 Step 3: Calculate the impact ratio

The **impact ratio** compares each group's selection rate with the highest selection rate.

Formula:
Impact ratio = Group selection rate ÷ Highest selection rate
Group A:
0.50 ÷ 0.50 = 1.00
Group B:
0.35 ÷ 0.50 = 0.70
Group C:
0.45 ÷ 0.50 = 0.90

6.11.6 Step 4: Apply the four-fifths rule

The four-fifths rule says that if a group's selection rate is less than 80% of the highest group's rate, that result may be evidence of adverse impact and should be investigated. The EEOC describes this as a general rule of thumb, not an automatic legal conclusion. (EEOC)

Group B's impact ratio is 0.70, which is below 0.80. That is a warning sign.

6.11.7 Step 5: Interpret the result

A warning sign does not prove the AI tool is illegal or discriminatory. It means HR should investigate before trusting the tool. The team should ask:

Was the AI recommendation actually job-related? What features did the tool use? Did it rely on proxies such as school prestige, employment gaps, ZIP code, or résumé style? Were applicants with disabilities disadvantaged by the format? Were candidates able to request accommodation? Were humans reviewing the recommendations carefully? Would a different screening method produce less adverse impact while still identifying qualified candidates?

6.11.8 Spreadsheet Version

Create columns named:

Group Applicants Recommended Selection Rate Impact Ratio Flag
In the Selection Rate column, use:
=C2/B2
In the Impact Ratio column, use:
=D2/MAX($D$2:$D$4)
In the Flag column, use:
=IF(E2<0.8,"Review needed","No 4/5ths flag")
Then copy the formulas down for all groups.

6.11.9 Lab Reflection

The most important result is not the number itself. The important result is the conversation the number forces HR to have. A responsible organization does not say, "The AI ranked them, so we are done." It says, "Can we justify this process? Can we explain it? Can we test it? Can people challenge it? Does it actually help us hire qualified people fairly?"

That is the difference between using AI as a tool and hiding behind AI as an excuse.

6.12 Final Takeaways

AI tools are definitely being used in Human Resources. The biggest uses are in recruiting, employee support, payroll assistance, learning, performance review drafting, workforce analytics, scheduling, and monitoring. The most mature deployments are often embedded in large HR software platforms rather than built from scratch by HR departments.

The benefits are real but practical: faster communication, less repetitive administration, better search across documents, improved scheduling, more consistent workflows, and stronger analytics. The risks are also real: biased screening, privacy invasion, inaccurate recommendations, over-monitoring, weak vendor transparency, and decisions that humans approve without meaningful review.

The best way to understand HR AI is not to ask, "Will AI replace HR?" A better question is: **Which HR tasks are being automated, who is affected, what data is being used, how are errors caught, and who remains accountable?**

In responsible organizations, AI helps HR professionals serve people better. In irresponsible organizations, AI can make unfair or careless systems faster and harder to challenge. The technology matters, but the governance matters just as much.

7. AI Tools in Financial Forecasting and Budgeting

Course Outcome

VCCS-7. Financial Forecasting and Budgeting Using AI: Apply AI-driven financial forecasting and budgeting techniques to predict trends, manage risks, and optimize business financial strategies.

AI tools are definitely being used in financial forecasting and budgeting in 2026. They are not usually "robot CFOs" making final budget decisions on their own. In most organizations, they are **assistive tools** built into spreadsheets, accounting software, enterprise planning platforms, business intelligence dashboards, and enterprise resource planning systems.

The best current evidence shows a mixed picture: adoption is real, but uneven. The Association for Financial Professionals reported in its 2025 FP&A Benchmarking Survey that **23% of finance planning and analysis professionals were using AI tools on a daily, weekly, or monthly basis**, while another **40% were testing AI and planned to implement it within the next year**. The same survey found that spreadsheets are still everywhere: **96% of respondents used spreadsheets for planning**, and **93% used them for reporting** at least daily or weekly. It also found that unreliable and inaccessible data were major barriers. (AFP)

Another 2025 FP&A Trends survey found a more cautious picture: only **8% of surveyed FP&A teams were actively using machine learning or predictive analytics in forecasting**, while **18% were using generative AI in some capacity**. Among the generative AI users, most were using it for communication and decision support; only a smaller share used it directly for forecasting or scenario modeling.

So the answer is: **yes, AI is being used in financial forecasting and budgeting, but mostly as a co-pilot, not an autopilot.** It helps analysts clean data, identify patterns, generate baseline forecasts, explain variances, test scenarios, and draft reports. Humans still approve budgets, challenge assumptions, decide tradeoffs, and take responsibility for the numbers.

7.1 What financial forecasting and budgeting mean

Before we talk about AI, we need to define the finance work.

Financial forecasting means estimating what will happen in the future: next month's sales, next quarter's cash flow, next year's payroll expense, or how much inventory the company will need. A forecast is not a promise. It is an informed estimate based on current data, historical patterns, and assumptions about the future.

Budgeting means deciding how much money the organization plans to spend and earn over a future period. A budget is usually more formal than a forecast. It can set targets for departments, managers, hiring, marketing, equipment purchases, debt repayment, and profit.

FP&A, short for **financial planning and analysis**, is the business function that prepares forecasts, budgets, performance reports, and "what if?" analyses for management. FP&A teams help leaders answer questions such as:

"How much cash will we have in 13 weeks?"

"What happens if sales are 5% lower than expected?"

"Can we afford to hire 20 more people?"

"Why did actual expenses come in higher than budget?"

For decades, this work has relied heavily on spreadsheets. Spreadsheets are flexible and familiar, but they also create problems. Different departments may use different versions of a file. Manual copy-and-paste work can introduce errors. Analysts can spend more time collecting and checking numbers than actually thinking about the business.

That is where AI tools are starting to matter.

7.2 The main types of AI being used

The phrase **artificial intelligence**, or **AI**, refers to computer systems that perform tasks that normally require human-like judgment, pattern recognition, language processing, or prediction. In finance, the most relevant forms are **machine learning**, **predictive analytics**, and **generative AI**.

Machine learning, or **ML**, is a type of AI that learns patterns from examples. In forecasting, a machine learning model might study past sales, prices, holidays, weather, website traffic, and customer behavior, then estimate future sales.

Predictive analytics means using data to estimate what is likely to happen next. Not all predictive analytics is advanced AI. Some forecasting tools use older statistical techniques, such as moving averages or exponential smoothing, which are mathematical methods for smoothing out noisy historical data. These are still useful and often more trustworthy than more complex models.

Generative AI is AI that creates new text, code, images, summaries, or explanations from prompts. A prompt is the instruction a user types or speaks to an AI tool. In finance, generative AI is commonly used to summarize results, draft budget narratives, answer questions about a report, or explain why actual results differed from forecast.

A fourth category is **automation**, which means software completing repeated tasks with less manual effort. Automation is not always AI. For example, automatically copying sales data from an accounting system into a planning model may be automation, even if no machine learning is involved. In real business systems, automation and AI are often combined.

7.3 Where AI fits in the forecasting and budgeting workflow

A typical financial forecasting and budgeting process has several stages. AI tools can assist at almost every stage.

7.3.1 Stage 1: Collecting and preparing data

A finance team may need data from an **ERP system**, short for **enterprise resource planning** system. An ERP is the main software system many organizations use to manage accounting, purchasing, inventory, payroll, invoices, and financial records. The team may also need data from a **CRM system**, short for **customer relationship management** system, which tracks sales opportunities and customer activity. Other data might come from payroll systems, bank accounts, point-of-sale systems, websites, inventory platforms, and spreadsheets.

This first stage is often the hardest. AI does not magically fix messy data. If product names are inconsistent, customer records are duplicated, sales regions are coded differently across systems, or managers use different spreadsheet formats, the model will struggle.

This is not a minor issue. The AFP survey found that many FP&A professionals face problems with data reliability, data accessibility, legacy systems, and disconnected tools. (AFP) FP&A Trends also reported that many teams are still working with older systems, and that spreadsheet-heavy teams spend a large share of their time on data collection and validation rather than insight and action.

AI tools can help by identifying missing values, detecting unusual entries, suggesting data transformations, matching similar categories, or converting messy tables into cleaner formats. Microsoft, for example, says its Finance capabilities in Microsoft 365 Copilot can connect to systems of record such as Dynamics 365 Finance or SAP and help users analyze forecast variances, spot anomalies, and prepare finance workflows inside familiar tools such as Excel and Outlook. (Microsoft) Microsoft also says Copilot in Excel can use Python-backed analysis for tasks such as forecasting and risk analysis from natural-language instructions. (Microsoft Support)

The important point: AI forecasting starts with data plumbing. Clean data is not glamorous, but it is the foundation.

7.3.2 Stage 2: Creating a baseline forecast

A **baseline forecast** is the model's first estimate of what will happen if current trends continue. For example, it might estimate next quarter's revenue based on last year's revenue, recent growth, seasonality, and known sales patterns.

Seasonality means a regular pattern that repeats at certain times. A retailer may sell more in December. A landscaping company may earn more in spring and summer. A college bookstore may see spikes at the beginning of each semester.

Forecasting tools can use many methods. Some are simple, such as:

A **moving average**, which averages recent periods to smooth out ups and downs.

Exponential smoothing, which gives more weight to recent data than older data.

ARIMA, a statistical forecasting method often used for time-series data, meaning data recorded over regular time intervals.

More advanced tools may use machine learning models that consider many variables at once. For example, revenue may depend on sales pipeline, price changes, customer churn, advertising spend, hiring levels, subscription renewals, weather, commodity prices, or macroeconomic conditions.

Oracle's Enterprise Performance Management tools include predictive planning features that use historical data to create forecasts, compare predictions with existing plans, and produce prediction ranges such as best-case and worst-case outcomes. Oracle's documentation describes forecasting methods such as moving averages, exponential smoothing, and ARIMA, and explains that the software can compare methods using error measures. (Oracle Docs) Oracle's Intelligent Performance Management features also look for forecast variance, forecast bias, anomalies, trends, and correlations in planning data. (Oracle Docs)

Anaplan's PlanIQ is another example. Anaplan says PlanIQ combines statistical forecasting, AI, and machine learning to generate predictions from internal and external data, compare model performance, and integrate forecasts into planning workflows. (Anaplan Inc)

7.3.3 Stage 3: Connecting forecasts to business drivers

A **business driver** is a factor that directly affects financial results. For a software company, drivers might include number of subscribers, renewal rate, average subscription price, and cloud hosting cost. For a restaurant, drivers might include customer traffic, average ticket size, food cost, labor hours, and rent. For a hospital, drivers might include patient volume, staffing ratios, reimbursement rates, and supply costs.

AI becomes more useful when it connects numbers to drivers. A forecast that says "revenue will be $12 million" is less helpful than one that says, "Revenue is likely to be $12 million because renewal rates improved, new bookings increased, and churn declined."

Workday describes AI features in Workday Adaptive Planning that support variance analysis, predictive forecasting, what-if scenarios, and conversational questions against financial plans. Workday says its tools can help users identify forecast drivers, compare scenarios, and track forecast confidence. (Workday)

This is especially important in budgeting. A budget is not just a prediction. It is a choice. Leaders may ask: Should we spend more on marketing? Should we delay hiring? Should we raise prices? AI can estimate possible outcomes, but managers must decide which tradeoffs fit the organization's goals and risks.

7.3.4 Stage 4: Scenario planning

Scenario planning means building multiple versions of the future. Instead of asking for one answer, the team asks:

"What if revenue grows 10%?"
"What if revenue falls 5%?"
"What if wages rise?"
"What if a major customer leaves?"
"What if we open a new location?"

AI tools can make scenario planning faster by adjusting many related assumptions at once. For example, if sales increase, the model may also increase inventory purchases, payment processing fees, warehouse labor, sales commissions, and taxes. If hiring is delayed, the model may reduce payroll costs but also reduce expected production capacity or customer service levels.

This is one of the most useful real-world applications because businesses rarely need a single perfect forecast. They need a range of possible outcomes and a plan for responding.

Coca-Cola HBC, a large bottling partner in the Coca-Cola system, describes using machine learning algorithms in demand forecasting and planning tools that support scenario planning, production sequencing, and medium- to long-term planning. This is not just finance department budgeting; it is operational planning that feeds financial budgets through expected demand, production needs, and inventory decisions. (CCH Group Website)

7.3.5 Stage 5: Variance analysis

Variance analysis means comparing actual results with the budget or forecast and explaining the difference. A variance is simply a difference between expected and actual results.

Suppose a department budgeted $100,000 for travel but spent $130,000. The variance is $30,000 unfavorable because the department spent more than planned. The next question is why. Was there a major conference? Higher airfare? More customer visits? Poor approval controls? A one-time event?

AI tools can scan transactions and highlight unusual changes. They can identify departments, vendors, products, regions, or accounts that changed more than expected. They can also draft explanations for finance teams to review.

Microsoft's Finance capabilities in Copilot include variance analysis features that identify anomalies and shifts, explain drivers, and draft management reporting narratives. (Microsoft) NetSuite Planning and Budgeting also describes AI and machine learning features that analyze multiple business drivers, surface insights, detect anomalies, and generate commentary explaining financial changes. (NetSuite)

This does not remove the analyst. It changes the analyst's work. Instead of spending hours finding the biggest variances, the analyst can spend more time asking whether the explanation makes business sense.

7.3.6 Stage 6: Narrative reporting

Finance teams do not only produce numbers. They also produce explanations for managers, boards, banks, investors, grant funders, and department heads.

Generative AI is useful here because it can turn a table into a readable first draft. For example:

"Revenue was 4% above forecast due to higher renewal rates in the Northeast region. Payroll expense was 3% above budget because of overtime in customer support. Management should watch support staffing levels in the next quarter."

Business intelligence tools are adding this type of capability. In Power BI, Microsoft describes Copilot-supported narrative visuals that create text summaries of report pages or visuals, while also warning users to review summaries for accuracy and understand that summaries depend on the data shown in the report. (Microsoft Learn)

That warning matters. A generative AI summary can sound confident even when it is incomplete or wrong. A good finance workflow treats AI-written narratives as drafts that must be checked against the numbers.

7.4 Examples from large, mid-sized, and small organizations

AI forecasting and budgeting looks different depending on the size of the organization.

7.4.1 Large enterprises

Large companies often use dedicated **EPM systems**, short for **enterprise performance management** systems. EPM software supports planning, budgeting, forecasting, consolidation, reporting, and performance management across departments and business units.

A large company may have thousands of cost centers, multiple currencies, international tax issues, complex supply chains, and strict access controls. In that setting, AI forecasting is usually embedded in a governed platform, not a random chatbot.

A current example comes from Arm Holdings, the semiconductor design company. MIT Sloan reported in 2026 that Arm's finance team had moved from forecasting royalties with Excel and other tools to using AI to build forecasts. Arm's CFO explained that royalty forecasting is important because a large share of Arm's revenue comes from royalties paid on billions of chips per quarter. (MIT Sloan)

This example is useful because it shows AI being used where the finance problem is data-heavy and recurring. The company still needs finance judgment, but AI can help process patterns at a scale that would be difficult in manual spreadsheets.

7.4.2 Mid-sized companies

Mid-sized companies often use platforms such as NetSuite, Workday Adaptive Planning, Anaplan, OneStream, Planful, Pigment, or similar planning tools. They may not have huge data science teams, but they still need better forecasting than manual spreadsheet rollups.

Workday has published a customer story about Team Car Care, the largest Jiffy Lube franchisee, moving from dozens of spreadsheets to a centralized planning platform and using Workday AI-related features for forecasting and planning. Because this is vendor-published, it should be read as a customer testimonial rather than independent proof of ROI, but it does show the type of deployment companies are buying: replacing spreadsheet-heavy planning with a centralized platform that includes predictive tools. (Workday)

Anaplan has also published customer material describing how The Coca-Cola Company uses Anaplan to work with bottling partners on revenue growth management planning. Again, because this is vendor-

published, it should be treated as evidence of product use and customer positioning, not neutral measurement of financial impact. (Anaplan Inc)

7.4.3 Small businesses

Small businesses usually do not start with enterprise AI platforms. They start with accounting software, bank feeds, spreadsheets, and maybe a payroll or point-of-sale system.

For a small business, the most important forecast may be **cash flow**, which means the timing of money coming in and going out. A company can be profitable on paper but still run short of cash if customers pay late, inventory must be purchased early, or payroll comes due before revenue is collected.

Intuit's QuickBooks documentation describes AI-assisted cash flow forecasts in Intuit Enterprise Suite that use roughly 18 months to two years of past financial data to generate a 13-week forecast, with monthly views up to 12 months. (QuickBooks) Xero's cash flow tools also show how small-business forecasting often uses invoices, bills, expected payment dates, recurring transactions, and user-created scenarios, even when the tool is not necessarily presented as advanced AI. (Xero)

For small businesses, the practical value is not "AI strategy." It is answering questions such as:

"Can I make payroll next month?"

"Can I buy new equipment?"

"What happens if my largest customer pays two weeks late?"

"Do I need a line of credit before the slow season?"

7.5 What AI is good at in finance planning

AI tools are strongest when the task has clear data, repeated patterns, and a need for speed.

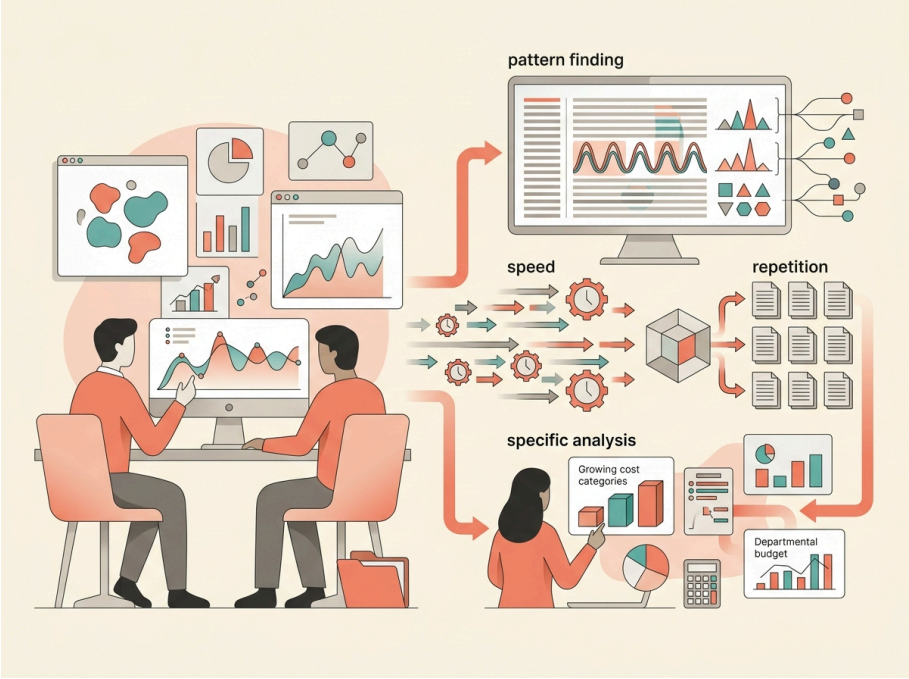

They are good at finding patterns across many rows of historical data. They can spot recurring seasonality, unusual spikes, delayed payments, expense categories that are growing quickly, or departments that consistently over-budget or under-budget.

They are good at producing a first forecast quickly. A finance analyst can then compare that forecast with management's judgment, sales pipeline information, and known business events.

They are good at scenario modeling. Instead of rebuilding a spreadsheet manually, an analyst can ask the system to test a lower-sales case, a higher-cost case, or a delayed-hiring case.

They are good at summarizing. A tool can draft a plain-English explanation of budget variances, highlight the biggest changes, or prepare a first version of a management report.

They are also useful for democratizing access to finance data. **Democratizing access** means making data easier for non-specialists to use. A department manager may not know how to build a pivot table or write a database query, but they can ask, "Why is my travel budget over plan?" A governed AI assistant can help them find the answer without waiting for an analyst to manually prepare a custom report.

7.6 What AI is not good at

AI forecasting is not magic. It has several serious limits.

First, AI depends on data quality. If the historical data is incomplete, mislabeled, biased, duplicated, or outdated, the forecast may be wrong. MIT Sloan's 2026 discussion of AI in finance emphasized that finance teams need clean data and that analytics projects often require substantial work just to acquire and clean the data. (MIT Sloan)

Second, prediction is not the same as judgment. A model might predict that sales will grow next quarter, but it cannot automatically know whether the company should hire more staff, raise prices, borrow money, or enter a new market. Those are business decisions involving risk, ethics, strategy, and human consequences.

Third, forecasting is not the same as causation. **Causation** means one thing actually causes another thing to happen. A model may find that marketing spend and revenue move together, but that does not prove that spending more on marketing will cause revenue to rise. Research on machine learning for FP&A has emphasized that prediction tools are useful, but planning and resource allocation often require causal reasoning, not just pattern recognition. (Springer)

Fourth, generative AI can make mistakes. A **large language model**, or **LLM**, is a generative AI system trained to predict and produce text. LLMs are powerful writing and summarizing tools, but they are probabilistic, meaning they generate likely answers rather than guaranteed correct answers. MIT Sloan quoted finance leaders warning that LLMs should be treated as a first draft or assistant, not the final answer in finance. (MIT Sloan)

Fifth, models can drift. **Model drift** happens when the real world changes and the old data no longer predicts the future well. A restaurant's 2023 sales pattern might not work after a new competitor opens nearby. A retailer's pandemic-era e-commerce pattern might not match later customer behavior. A company's cost model may break when interest rates, wages, exchange rates, or supply chain costs change.

Sixth, security and access control matter. Financial forecasts may include payroll plans, layoffs, acquisition discussions, pricing strategy, cash problems, or investor-sensitive information. AI tools used in finance need permissions, audit trails, role-based access, and data governance. Microsoft and Workday both emphasize governed access and enterprise security in their finance AI offerings. (Microsoft)

7.7 How the finance job changes

AI does not remove the need for finance professionals. It changes what good finance work looks like.

A traditional junior analyst might spend much of the week downloading data, fixing spreadsheet formats, updating formulas, checking totals, and preparing slides. With AI-enabled tools, some of that work can be automated or accelerated.

But the analyst still needs to ask better questions:

"Is this forecast based on the right driver?"

"Did the model overreact to one unusual month?"

"Does the forecast match what the sales team is seeing?"

"Are we confusing correlation with causation?"

"Would this budget decision hurt customer service or employee retention?"

"Can managers understand and trust this explanation?"

The future finance professional needs both technical fluency and business judgment. Technical fluency does not mean every finance worker must become a data scientist. It means they should understand enough about data, models, prompts, controls, and limitations to use AI responsibly.

7.8 Hands-on Lab: Build and Check an AI-Assisted Cash Forecast

7.8.1 Lab goal

In this lab, you will act as a junior FP&A analyst for a small café. Your job is to create a simple three-month cash forecast, then use an AI assistant as a reviewer—not as the final authority.

You can complete this lab in Excel, Google Sheets, or another spreadsheet tool. You may also use a generative AI tool if your instructor allows it.

7.8.2 Dataset

Copy this table into a spreadsheet.

Month	Revenue	Payroll	Rent & Utilities	Inventory Purchases	Other Expenses
Jan 2025	42,500	15,500	6,200	12,200	4,100
Feb 2025	43,800	15,500	6,200	12,400	4,200
Mar 2025	46,000	15,800	6,200	13,100	4,000
Apr 2025	48,500	16,000	6,300	13,900	4,300
May 2025	51,200	16,200	6,300	14,700	4,500
Jun 2025	54,000	16,500	6,300	15,400	4,800
Jul 2025	57,500	16,800	6,400	16,200	5,100
Aug 2025	56,900	16,800	6,400	16,000	5,000
Sep 2025	52,300	16,600	6,400	14,900	4,700
Oct 2025	49,700	16,400	6,500	14,100	4,500
Nov 2025	50,800	16,600	6,500	14,300	4,600
Dec 2025	59,000	17,200	6,600	16,900	5,300
Jan 2026	47,400	17,300	6,700	13,300	4,300
Feb 2026	48,900	17,300	6,700	13,700	4,400
Mar 2026	51,600	17,600	6,700	14,600	4,400
Apr 2026	53,900	17,800	6,800	15,200	4,500
May 2026	56,200	18,000	6,800	15,900	4,700
Jun 2026	58,000	18,100	6,800	16,300	4,900

7.8.3 Step 1: Calculate net cash flow

Add a new column called **Net Cash Flow**.
Use this formula:
Revenue - Payroll - Rent & Utilities - Inventory Purchases - Other Expenses
For June 2026, the calculation is:
58,000 - 18,100 - 6,800 - 16,300 - 4,900 = 11,900
That means the café generated an estimated positive cash flow of $11,900 in June 2026.

7.8.4 Step 2: Create a simple baseline revenue forecast

Use a three-month moving average for July 2026 revenue.
Average April, May, and June 2026 revenue:
(53,900 + 56,200 + 58,000) ÷ 3 = 56,033
So a simple baseline forecast for July 2026 revenue is **$56,033**.
This is intentionally simple. It does not fully account for seasonality, but it gives you a starting point.

7.8.5 Step 3: Add a seasonal adjustment

Now look at summer 2025. Revenue rose from June 2025 to July 2025:
57,500 - 54,000 = 3,500
That was an increase of about 6.5%.
Apply a cautious seasonal adjustment to June 2026 revenue:
58,000 × 1.065 = 61,770
Now you have two possible July revenue estimates:
Simple moving average: **$56,033**
Seasonally adjusted estimate: **$61,770**

This shows an important forecasting lesson: the method matters. Different reasonable methods can produce different results.

7.8.6 Step 4: Build three scenarios

Create three July 2026 scenarios.

Scenario	Revenue Assumption
Low case	$56,033
Base case	$61,770
High case	$64,859

The high case is 5% above the base case:
61,770 × 1.05 = 64,859
Now estimate expenses. Use these assumptions:
Payroll: $18,300 Rent & utilities: $6,900 Inventory purchases: 28% of revenue Other expenses: $5,000
For the base case, inventory would be:
61,770 × 0.28 = 17,296
Base-case net cash flow would be:
61,770 - 18,300 - 6,900 - 17,296 - 5,000 = 14,274

7.8.7 Step 5: Use an AI assistant carefully

Paste the data and your formulas into an AI tool and use this prompt:
You are helping an FP&A analyst review a small café cash forecast. Use only the data provided. Create a July 2026 low, base, and high cash-flow forecast. State your assumptions clearly. Do not invent new data. Show the arithmetic. Flag any risks or missing information.
Now compare the AI output with your spreadsheet. Do not assume the AI is correct.
Check these items:
Did it use only the data you provided?
Did it calculate inventory correctly as 28% of revenue?
Did it subtract all expense categories?
Did it label assumptions clearly?
Did it make up facts, such as customer counts or menu prices?
Did it explain uncertainty?
This is how AI should be used in finance: as a reviewer, accelerator, and idea generator—not as an unchecked authority.

7.8.8 Step 6: Write a short finance comment

Write a three-sentence management comment:
"The July base-case forecast shows positive cash flow of about $14,274, mainly because revenue is expected to improve during the summer season. The largest risk is that the seasonal increase may not repeat, which would bring cash flow closer to the low case. Management should monitor weekly sales and inventory purchases before committing to major new spending."
That comment is simple, but it does what finance communication should do: it connects the numbers to the business decision.

7.8.9 Key takeaway

AI tools are already being used in financial forecasting and budgeting, but the most common real-world pattern is **human-led, AI-assisted planning**. AI helps finance teams prepare data, create baseline forecasts, test scenarios, spot variances, and draft explanations. The strongest organizations do not simply "trust the AI." They combine AI outputs with clean data, business knowledge, internal controls, and human judgment.

8. How AI Tools Are Used in Business Process Automation and RPA

Course Outcome

VCCS-8. AI in Business/Designing Business Report Prompts: Examine the benefits of AI-powered process automation and robotic process automation (RPA) in improving business productivity and efficiency.

AI tools are already being used in **Business Process Automation** and **Robotic Process Automation** in businesses and government agencies. The important point is *how* they are being used. In 2026, most organizations are not handing entire business departments over to free-roaming artificial intelligence. Instead, they are combining older automation tools with newer AI tools inside controlled workflows: extracting information from documents, routing cases, drafting responses, summarizing tickets, checking policy rules, and asking a human to review exceptions.

That may sound less dramatic than the popular idea of "AI agents running the company," but it is more realistic and more important. Most business work is not one giant decision. It is a chain of smaller tasks: receive a form, read it, check it against a policy, update a system, notify someone, wait for approval, and keep a record. AI is being inserted into those chains where older automation was too rigid.

8.1 The basic vocabulary

A **business process** is a repeatable series of steps that produces a business result. Hiring an employee, paying an invoice, resetting a password, approving a refund, ordering supplies, or processing an insurance claim are all business processes.

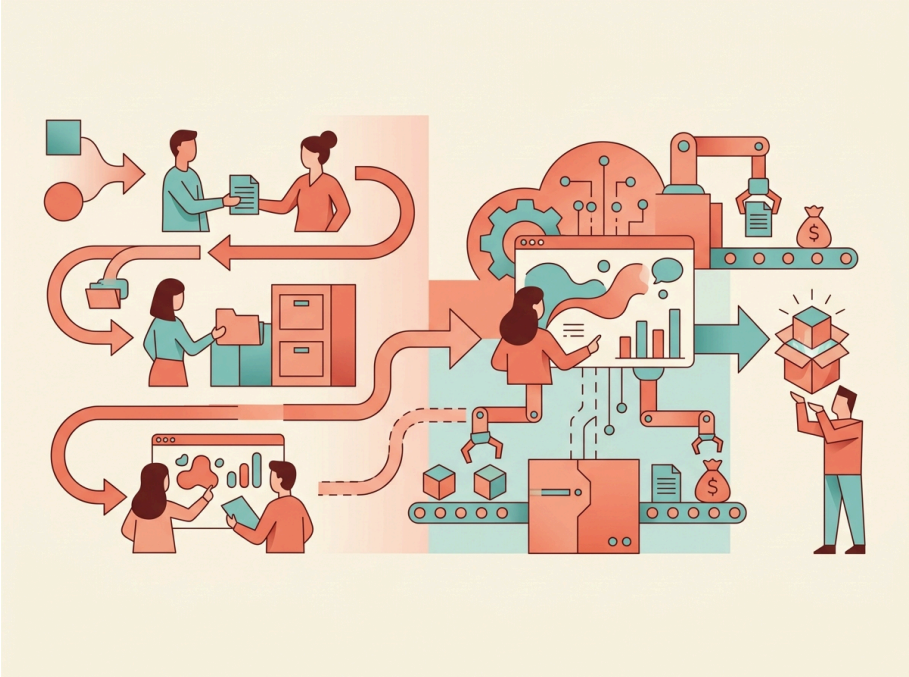

Business Process Automation, often shortened to **BPA**, means using software to automate all or part of a business process. BPA usually focuses on the process as a whole: who starts it, which systems are involved, what approvals are needed, what records are created, and how the process is measured.

Robotic Process Automation, or **RPA**, is a narrower kind of automation. An RPA "bot" is software that imitates human actions on a computer: clicking buttons, copying and pasting data, opening emails, downloading files, filling forms, and logging into business applications. The "robot" is not a physical robot. It is software running on a computer or virtual machine. A U.S. General Services Administration audit described RPA bots as tools that simulate human actions such as copying data, filling forms, signing into applications, and sending emails. The same audit also warned that bots can perform many actions quickly, so they need strong security controls.

An **application programming interface**, or **API**, is a software-to-software connection. If one system can send information directly to another system through an API, that is usually cleaner than having a bot click around a screen. RPA is especially useful when the older system does not have a good API, or when replacing the system would be too expensive.

Artificial intelligence, or **AI**, means software that performs tasks that normally require pattern recognition, prediction, language understanding, or decision support. In BPA and RPA, AI commonly helps software read documents, classify requests, summarize text, detect anomalies, recommend a next step, or generate a draft response.

Generative AI is AI that creates new text, images, code, or other content. A **large language model**, or **LLM**, is a generative AI system trained to work with human language. In automation, LLMs are often used to summarize a customer complaint, draft an email, turn a plain-English request into a workflow, or help an employee find the right policy.

An **AI agent** is software that uses AI to pursue a goal by choosing steps and using tools. In business automation, that might mean an agent that reads an employee request, decides whether it is about payroll or benefits, opens the right workflow, drafts a reply, and escalates the case if it is risky. The key phrase is "bounded workflow." Serious organizations usually restrict agents to approved tools, approved data, and human review points.

8.2 From rules-based automation to AI-enabled automation

Older automation was mostly **rules-based**. A rule might say: "If the invoice is under $500 and the vendor is approved, send it to the department manager." Rules-based automation is still extremely useful because it is predictable. The same input should produce the same output.

The problem is that business work often arrives in messy human form. A customer writes a paragraph instead of selecting a category. A supplier sends a PDF invoice in a different format. An employee asks, "Can I take leave next Friday?" instead of filling out the correct HR form. Older automation struggled with this kind of unstructured information.

That is where AI fits. AI can turn messy information into structured information. For example, it can read a PDF invoice and extract the vendor name, invoice number, amount, due date, and purchase order number. It can read a support ticket and classify it as "password reset," "software access," or "network issue." It can summarize a long email thread so a human reviewer can act faster.

UiPath, one of the major RPA vendors, describes its current platform as combining AI agents, user-interface automation, API integration, document understanding, process mining, task mining, and governance tools. Its 2026 annual filing says its platform uses AI technologies including machine learning, natural language processing, and computer vision to support agents and automations. (UiPath, Inc.)

That description shows the larger industry shift: RPA is no longer only about screen-clicking bots. It is becoming part of a broader automation stack that includes AI, workflow orchestration, document processing, and monitoring.

8.3 Where AI is actually used in BPA and RPA

One of the biggest uses is **intelligent document processing**. This means using AI to read documents and convert them into structured data. Common examples include invoices, receipts, purchase orders, loan documents, claims forms, tax forms, contracts, shipping documents, and benefit applications.

A typical invoice automation process might work like this. An invoice arrives by email. AI extracts the vendor, amount, invoice number, date, and purchase order. A rule checks whether the purchase order matches company records. If the amount is small and everything matches, the system routes it for normal approval or posts it to the accounting system. If the amount is large, the vendor is new, or the numbers do not match, the system sends it to a human for review.

Microsoft's AI Builder documentation describes document processing models that can be trained and then used in Power Automate or Power Apps. This is a good example of how AI document extraction is being connected directly to workflow automation rather than used as a standalone tool. (Microsoft Learn)

Another major use is **ticket and case handling**. A ticket is a request that needs to be tracked, such as "my laptop will not connect to Wi-Fi" or "I need access to the finance folder." A case is a broader service record, such as an HR issue, customer complaint, or legal review request. AI can classify the ticket, summarize the issue, suggest a response, recommend a knowledge-base article, or trigger an automated workflow.

ServiceNow, a major enterprise workflow platform, says its AI platform supports AI agents that can trigger IT provisioning, payroll setup, compliance checks, and facilities access, while using human oversight and guardrails. Its 2025 annual filing also describes AI support for incident triage, summaries, and resolution recommendations in IT service management. (SEC) (SEC)

AI is also used in **HR automation**. HR stands for human resources, the department that handles employee-related processes such as hiring, onboarding, benefits, leave, transfers, and offboarding. ServiceNow's filing describes HR service delivery tools that automate routine HR tasks and support onboarding, leave, transfers, offboarding, and service requests. (SEC)

In procurement, which means buying goods and services for an organization, AI can help create purchasing requests, check spending policies, suggest categories, and route approvals. ServiceNow describes source-to-pay AI agents that can initiate procurement requests, prefill details, and check spending policy. (SEC)

Legal and contract work is another area. AI can summarize a contract, identify nonstandard language, recommend clauses, or route the document to legal staff. This does not mean the AI becomes the lawyer. It means the AI helps screen, summarize, and prepare work for a qualified human reviewer. ServiceNow's filing describes AI tools for legal and contract work that detect nonstandard language, recommend clauses, and generate summaries. (SEC)

8.4 Process mining and task mining: finding what to automate

Before a company automates a process, it needs to understand the process. This is harder than it sounds. The official process diagram might say that invoices go from "received" to "approved" to "paid." The real process might involve five email threads, three spreadsheets, a manager who is always late, and a clerk who manually fixes vendor names every Friday.

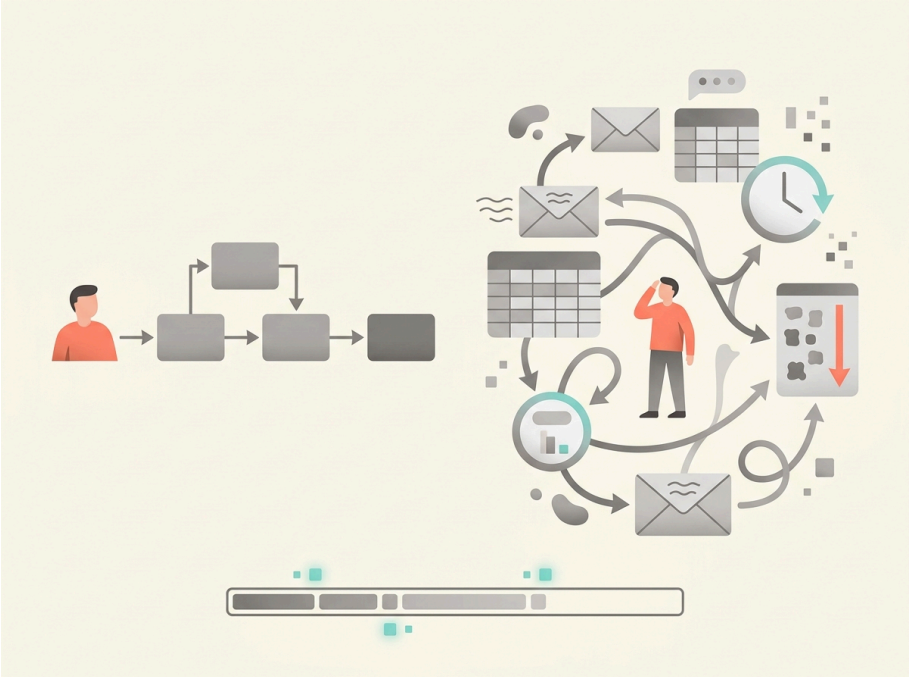

Process mining uses event logs from business systems to reconstruct how a process actually works. An event log is a record of what happened in a system: when a case was opened, who approved it, when it changed status, and when it closed. Process mining can show bottlenecks, rework, delays, and variations.

Task mining focuses on the desktop work people perform, such as copying data from one application into another. It can help identify repetitive work that might be automated. Because task mining may observe employee computer activity, it raises privacy and labor concerns. Organizations should be transparent about what is being collected and why.

UiPath's annual filing describes process mining and task mining as part of its automation platform. Microsoft's Power Automate materials also describe process and task mining as tools for understanding and optimizing processes before automating them. (UiPath, Inc.) (Microsoft Learn)

This is an important point for students: automation should not begin with the question, "Where can we use AI?" A better question is, "Where is the process slow, repetitive, error-prone, expensive, or frustrating?" AI is only useful when it improves a real process.

8.5 RPA bots, AI agents, and workflows are different tools

It is easy to confuse bots, agents, and workflows, so let's separate them.

A **workflow** is the organized sequence of steps in a process. For example: receive request, classify request, ask for approval, update record, notify employee.

An **RPA bot** performs specific computer actions, usually through a user interface. It might open a website, enter data, download a report, and upload that report into another system.

An **AI agent** interprets information and chooses actions within limits. It might decide whether a request is about payroll, benefits, or facilities. It might choose which workflow to start. It might draft a response to a customer. But in responsible deployments, the agent is usually surrounded by rules, permissions, logs, and human review.

Microsoft's 2025 annual report says Power Automate is part of its business applications portfolio, along with Dynamics 365 and Power Apps. The same report says Microsoft has seen large-scale use of Copilot and Copilot Studio, with more than 230,000 organizations using Copilot Studio to extend Microsoft 365 Copilot or build agents with low-code and no-code tools. (Microsoft) (Microsoft)

Microsoft's Copilot Studio documentation also distinguishes deterministic flows from agent behavior. In plain English, a deterministic workflow is one where the same input should lead to the same result. That matters because many business processes, especially finance, HR, and compliance processes, need consistency and auditability. (Microsoft Learn)

The practical lesson is that companies are not simply replacing workflows with AI agents. They are often putting agents *inside* workflows.

8.6 What real deployments look like

A realistic AI-enabled automation often looks like this:

1. A request or document enters the system.
2. AI extracts or summarizes the information.
3. Rules check policy, amount, identity, or risk.
4. The workflow routes the item to the right person or system.
5. RPA or an API updates another application.
6. A human reviews exceptions or high-risk decisions.
7. The system logs what happened.

For example, consider employee onboarding. A new hire accepts an offer. The HR system creates an onboarding case. AI reads the job title and location. Rules decide which laptop, software, badge access, payroll setup, and training are needed. An RPA bot might enter data into an older facilities system. An API might create accounts in newer cloud systems. A manager approves unusual access. The employee receives automated instructions. The entire process leaves an audit trail.

Or consider customer support. A customer writes, "I was charged twice and the refund still has not arrived." AI classifies the message as a billing issue, summarizes the complaint, detects frustration, and pulls relevant account details. A workflow checks refund status. If the case is simple, the system drafts a response for an agent. If the amount is large or fraud is suspected, the case is escalated.

Or consider public benefits. A state agency receives thousands of forms. RPA can move data between systems, while AI can help read forms and classify documents. But public benefits are high-stakes because errors can affect people's food, housing, or medical care. A 2025 U.S. Department of Agriculture Food and Nutrition Service study of RPA in Supplemental Nutrition Assistance Program administration found that, as of January 2022, nine states were using RPA in SNAP. The study found some benefits, including lower payment error rates in one Georgia recertification process and cost benefits within one year for Georgia's RPA use, but it also found limits in measurement and did not find significant improvement on every processing-time measure. (USDA Food and Nutrition Service)

That balanced result is useful. Automation can help, but it does not magically fix every process.

8.7 Who is using these tools?

Large software companies are building AI into automation platforms. UiPath reports thousands of large customers and says its platform combines RPA, AI agents, document understanding, process mining, task mining, and governance. In its 2026 annual filing, UiPath reported $1.85 billion in annual recurring revenue and 2,565 customers with at least $100,000 in annual recurring revenue, suggesting that automation platforms are already embedded in many organizations. (UiPath, Inc.)

Microsoft is integrating AI automation into Power Automate, Power Apps, Dynamics 365, Microsoft 365 Copilot, and Copilot Studio. These products matter because many organizations already run Microsoft software for email, documents, collaboration, identity, and business applications. Microsoft's 2025 annual report says its business applications include Dynamics 365, Power Apps, and Power Automate, and it describes growth tied to AI-enabled tools and agentic scenarios. (Microsoft)

ServiceNow is embedding AI into IT, HR, customer service, procurement, security, legal, and operational workflows. Its annual filing emphasizes AI agents, human oversight, and enterprise workflow automation. (SEC)

SAP, a major enterprise resource planning vendor, is also adding AI to business processes. **Enterprise resource planning**, or **ERP**, means software that manages core business functions such as finance, procurement, inventory, supply chain, and human resources. SAP's own 2025 release materials describe AI features, Joule agents, and SAP Build tools for process automation and AI-driven workflow design. Because this is vendor-published material, it should be read as product documentation rather than independent proof of results. (SAP News Center) (SAP)

Government agencies are also using automation. The U.S. General Services Administration's Federal Automation Community of Practice says it includes more than 1,700 members from more than 100 federal departments and agencies, and that its 2025 inventory collected more than 3,000 automation use cases. (U.S. General Services Administration)

8.8 Why businesses use AI automation

The first reason is speed. A bot can move data or check records faster than a person doing repetitive keystrokes. AI can also reduce the time needed to read and summarize documents.

The second reason is consistency. A workflow can enforce the same routing rule every time. A human may forget a step at the end of a long day; software does not get tired. Of course, software can still be wrong if the rule is wrong, the data is wrong, or the AI model misreads the input.

The third reason is cost control. Organizations often have back-office work that grows with business volume: invoices, claims, tickets, account updates, compliance checks. Automation can help teams handle more work without hiring at the same rate. But students should be careful with dramatic cost-savings claims. Many vendor case studies report impressive percentages, but those numbers are often promotional unless the method is public.

The fourth reason is old-system integration. Many organizations still depend on legacy systems. A **legacy system** is an older technology that remains important even though it may be hard to modify. RPA can act as a bridge by using the same screens that human workers use.

The fifth reason is better measurement. A digital workflow can record when each step starts, where work waits, how many exceptions occur, and how often rework happens. That data can help managers improve the process.

8.9 The risks: bots can make mistakes faster than people

Automation risk is not imaginary. A bad manual process might create ten errors per day. A bad bot might create ten thousand errors before anyone notices.

The GSA Office of Inspector General audited GSA's RPA program and found that the agency needed stronger security controls. The audit noted that bots can perform many read, write, and delete actions quickly, and it raised concerns about access controls, monitoring, and decommissioned bots. It also noted that a prior review found GSA did not have enough evidence to support some claimed savings because it was not consistently verifying hours saved or tracking costs.

This is a perfect example of why automation governance matters. A bot is not "just a script." If it logs into systems, changes records, or moves data, it is a digital worker with permissions. It needs an owner, a purpose, a security review, access limits, monitoring, and a retirement plan.

AI adds more risks. A document model may extract the wrong amount from an invoice. A language model may summarize a complaint incorrectly. An agent may choose the wrong workflow. A chatbot may give an employee an outdated policy answer. These problems are especially serious in finance, healthcare, employment, public benefits, insurance, education, and legal settings.

The U.S. Office of Management and Budget's 2024 government AI guidance emphasizes AI governance, innovation, and risk management, including minimum practices for AI uses that can affect rights or safety. NIST's Generative AI Profile is designed to help organizations identify risks specific to generative AI and manage them across design, development, use, and evaluation. (NIST)

8.10 What good governance looks like

Good AI automation governance begins with an inventory. An **inventory** is a list of automations in use: what each one does, who owns it, which systems it touches, what data it uses, and what could go wrong.

Next comes access control. A bot should have only the permissions it needs. This is called **least privilege**. For example, a bot that reads invoice data should not also be able to change employee payroll records.

Organizations also need logs. A **log** is a record of what happened: what the bot did, when it did it, what data it used, and whether a human approved the action. Logs matter for debugging, security, compliance, and accountability.

Testing is essential. Before an automation goes live, teams should test normal cases, edge cases, and failure cases. An **edge case** is an unusual situation that may break the process, such as a negative invoice amount, a missing purchase order, a duplicate customer record, or a document in the wrong language.

Human review should be designed into the workflow. This is often called **human-in-the-loop** review. The point is not that humans must approve every tiny action. The point is that humans should review high-risk, uncertain, unusual, or irreversible actions.

Finally, organizations need change management. If a website layout changes, a screen-clicking bot may break. If a policy changes, an AI assistant may give old advice. If a model is updated, its behavior may change. Automation is not "set it and forget it." It is a living system.

8.11 What this means for jobs

AI automation changes work, but not always in the simple "robots replace people" way. In many offices, the first tasks automated are repetitive data movement, document sorting, basic classification, status updates, and standard replies. People still handle exceptions, judgment calls, customer relationships, policy interpretation, and process improvement.

New roles are also growing around automation. A **business analyst** studies how a process works and where it breaks down. An **RPA developer** builds and maintains bots. A **process owner** is responsible for the performance of a business process. A **data steward** helps ensure that data is accurate and well governed. An **AI governance specialist** helps review risk, compliance, privacy, and accountability. A **citizen developer** is a business user who builds simple apps or workflows with low-code tools, usually under IT supervision.

For BIT students, the opportunity is not only to learn coding. It is to learn how business processes, data, systems, people, and controls fit together. The valuable worker is the person who can say: "Here is the process, here is the bottleneck, here is what should be automated, here is what should not be automated, and here is how we will measure whether it worked."

8.12 Hands-on lab: design an AI-enabled invoice automation

This lab does not require paid enterprise software. You can complete it with a spreadsheet and any AI chatbot approved by your instructor. If you have access to Microsoft Power Automate Desktop, you can optionally build a small version of the workflow. Microsoft says Power Automate includes cloud flows, desktop flows, process and task mining, and AI-powered automation, but licensing varies; unattended desktop automation and premium features usually require paid plans. (Microsoft) (Microsoft)

8.12.1 Scenario

You work for a small company. Vendors send invoices by email. An accounting clerk opens each invoice, enters the information into a spreadsheet, checks whether the purchase order matches, and sends the invoice to a manager for approval. The company wants to automate the first review.

8.12.2 Sample data

Create a spreadsheet with these columns:

Invoice ID	Vendor	Amount	PO Match?	Description	Due Date
INV-1001	Northstar Office	245.80	Yes	Printer paper and toner	2026-06-15
INV-1002	GreenGrid Energy	3,850.00	Yes	Monthly electricity service	2026-06-20
INV-1003	Apex Consulting	1,200.00	No	Strategy workshop	2026-06-18
INV-1004	Metro Catering	480.00	Yes	Lunch for training event	2026-06-12
INV-1005	Unknown Vendor LLC	775.00	No	Rush service fee	2026-06-10

8.12.3 Step 1: Map the current process

Write the current manual process in plain English:

Receive invoice → open attachment → read invoice → enter data → check purchase order → decide approval path → email manager → update spreadsheet → file invoice.

Now mark each step as one of three types:

Rule-based automation: predictable steps such as checking whether the amount is over $1,000.

AI assistance: language or document tasks such as summarizing the invoice description or classifying risk.

Human review: approvals, exceptions, suspicious vendors, missing purchase orders, or high-dollar invoices.

8.12.4 Step 2: Create automation rules

Use these rules:

- If the amount is greater than $1,000, send to human review.
- If "PO Match?" is "No," send to human review.
- If the vendor name includes "Unknown," send to human review.
- If none of those are true, mark as "auto-approval candidate."

These rules are deterministic. They do not require AI.

8.12.5 Step 3: Add AI classification

Ask an AI tool to classify each invoice description as **Low Risk**, **Medium Risk**, or **High Risk**. Use a prompt like this:

You are helping review invoices. Classify each invoice as Low Risk, Medium Risk, or High Risk based only on the vendor name, amount, purchase order match, and description. Explain your reason in one sentence. Do not approve payment. Human reviewers make final decisions.

Add two new columns to the spreadsheet:

- AI Risk Level
- AI Reason

Then compare the AI classification with your deterministic rules. Did the AI notice anything your rules missed? Did it overreact? Did it make an assumption that was not supported by the data?

8.12.6 Step 4: Design the future process

Write the improved process:

Email receives invoice → document AI extracts fields → workflow checks rules → AI classifies risk → low-risk invoices become approval candidates → high-risk or uncertain invoices go to a human → approved invoices are posted to accounting system → all actions are logged.

In a real company, "posted to accounting system" might happen through an API. If no API exists, an RPA bot might enter the data into an older accounting application.

8.12.7 Step 5: Add governance controls

Add at least five controls:

- A human must approve all invoices over $1,000.
- A human must approve all invoices without a purchase order match.
- The bot can only access the invoice folder and accounting screen it needs.
- Every action is logged with date, invoice ID, and decision.
- The workflow stops and alerts a person if required data is missing.
- The AI classification is advisory, not final.
- The automation owner reviews errors weekly.

8.12.8 Step 6: Test the automation

Create three test cases:

1. A normal low-dollar invoice with a matching purchase order.
2. A high-dollar invoice with a matching purchase order.
3. A suspicious invoice from a new vendor with no purchase order.

For each test case, write the expected result. This is how real automation teams think: not just "Can we build it?" but "How will we know it works, and what happens when it fails?"

8.12.9 Lab deliverable

Submit a one-page automation brief with four sections:
Process: What process are you automating?
AI role: What does AI do, and what does it not do?
Automation logic: What rules decide the routing?
Controls: How do you prevent errors, fraud, privacy problems, or unsafe approvals?

8.13 The main takeaway

AI tools are definitely being used in BPA and RPA. The most common pattern is not "AI replaces the whole process." The real pattern is:

AI reads, classifies, summarizes, predicts, or drafts. Workflow software routes and records the work. RPA bots or APIs update systems. Humans review exceptions and high-risk decisions. Governance keeps the system accountable.

That combination is powerful because it matches how organizations actually operate. Businesses do not run on isolated AI tricks. They run on processes, systems, approvals, records, and responsibilities. AI becomes valuable when it helps those processes work faster, more accurately, and more transparently—without removing the controls that protect customers, employees, and the organization itself.

9. AI Tools in Product Design and Development

Course Outcome

VCCS-9. Product Design and Development: Discuss AI's impact on product development and innovation, including research, performance simulation, and accelerated innovation processes.

AI tools are already being used in product design and development, but not usually as "push a button and invent a product" machines. In 2026, the practical pattern is more ordinary and more important: AI is being added to the tools that product managers, designers, engineers, marketers, and software developers already use. It helps them summarize customer feedback, explore design alternatives, generate early mockups, write product requirements, produce code, run simulations, create test cases, and analyze post-launch data. The work is still human-led because product development involves tradeoffs: user needs, cost, safety, brand, manufacturability, legal risk, and timing.

9.1 What "product design and development" means

A **product** can be physical, digital, or both. A water bottle, a running shoe, a banking app, a medical device, a video game, and an electric vehicle dashboard are all products. **Product design** is the work of deciding what the product should be like: what problem it solves, who it serves, how it looks, how it feels, and how it behaves. **Product development** is the larger process of turning that idea into something real: researching users, writing requirements, designing prototypes, engineering the product, testing it, launching it, and improving it after release.

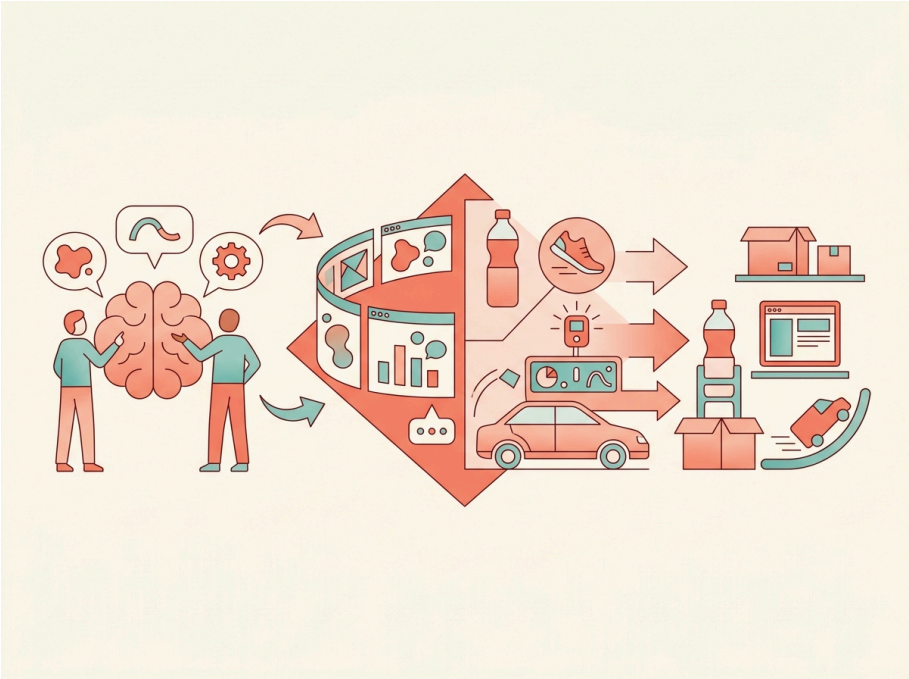

AI appears throughout that cycle. A 2025 academic review of AI in new product development found that AI is being applied across stages such as idea generation, design support, forecasting, recommendation, and optimization, but also warned that adoption is fragmented: many tools help one stage, while few connect the whole product lifecycle end to end. The same review notes that AI is often discussed too broadly, so useful analysis requires breaking product development into concrete phases and use cases. (Springer)

9.2 The AI terms you need first

Artificial intelligence, or **AI**, means computer systems that perform tasks associated with human intelligence, such as recognizing patterns, generating text, making predictions, or suggesting decisions. **Machine learning** is a common type of AI in which software improves by learning patterns from data rather than being programmed with every rule by hand. **Generative AI** is AI that creates new content, such as text, images, code, designs, audio, or 3D concepts. A **large language model**, or **LLM**, is a generative AI model trained on large amounts of text so it can respond to prompts, summarize information, draft documents, and help with code.

In product work, you will also hear **CAD**, which means **computer-aided design**: software used to create precise digital models of physical objects. **UX**, or **user experience**, is how a person experiences a product overall; **UI**, or **user interface**, is the screens, buttons, menus, and controls a user interacts with. A **prototype** is an early version of a product used for learning before full production. A **product requirements document**, or **PRD**, describes what a product or feature should do and why. A **digital twin** is a digital model of a real product, machine, factory, or process that can be used for simulation and "what if" testing before changing the real-world version.

9.3 AI in customer discovery: turning messy feedback into usable signals

The earliest product question is not "What can we build?" but "What problem is worth solving?" Product teams collect support tickets, app reviews, survey answers, sales notes, interview transcripts, call recordings, and usage data. AI tools help by summarizing long conversations, clustering similar complaints, detecting sentiment, and finding repeated themes. **Sentiment analysis** means using software to estimate whether text expresses positive, negative, or neutral attitudes.

For example, Productboard, a product-management software vendor, describes AI features that summarize customer feedback, help write feature briefs, and surface insights from long conversations. Its page also states that Productboard AI is powered by OpenAI and that customer data is not used to train models for other customers. Because this is vendor-published information, it should be treated as a description of product capability, not as independent proof of business value. (Productboard)

This use is attractive because product teams often drown in feedback. A product manager may have hundreds of comments like "the checkout page is confusing," "I couldn't find the return button," and "the app froze after payment." AI can group those into themes such as navigation problems, payment reliability, and unclear language. But AI does not automatically know which complaint matters most. A loud group of users may not represent the whole market. A summary may hide a crucial detail. Good teams therefore use AI as a sorting assistant, then verify findings against actual evidence.

9.4 AI in idea generation and concept design

Once a team understands a problem, it explores possible solutions. This is where generative AI became very visible. Designers can ask an AI tool to generate mood boards, packaging variations, interface sketches, product names, landing-page copy, visual styles, or early screen layouts. This helps teams create many options quickly, especially at the "rough draft" stage.

Adobe's 2025 annual report describes Firefly as a family of creative generative AI models spanning image, video, vector, audio, and more. It says Firefly is built into Adobe products and can create or edit images, videos, text effects, design templates, vector graphics, and audio through natural language prompts and reference assets. Adobe also describes Firefly Custom Models and Firefly Foundry, which let enterprises train or tune models around their own brand, product, or intellectual-property styles, and Firefly Services, which provides APIs for content generation, editing, and assembly. (Adobe)

This matters for product design because the early concept stage is visual and iterative. A team designing a new backpack, snack package, or mobile app onboarding flow can generate multiple directions before choosing one to refine. However, "more ideas" is not the same as "better ideas." Generative AI models learn from existing data, so they tend to remix familiar patterns. The 2025 product-development literature review warns that overreliance on historical data may favor incremental improvements over radical innovation. (Springer)

9.5 AI in UX/UI design and interactive prototypes

Digital product teams are using AI inside tools such as Figma, which is widely used for collaborative interface design. In 2026, Figma reported that its Q4 2025 net dollar retention rate rose to 136% "as Figma drove platform and AI adoption." The same report said weekly active users of Figma Make grew over 70% quarter over quarter, and that more than half of paid customers with over $100,000 in annual recurring revenue were building in Figma Make weekly during the three months ending December 31, 2025. (Figma Investor Relations)

Figma Make is an AI-assisted product that lets users start with a design and prompt their way to a functional prototype. Figma describes it as helping teams create high-fidelity prototypes, apply styling context from a design system, refine selected parts of a design through prompts, and connect to Supabase, a backend platform, to build apps with real data. (Figma)

This is a major shift in UX/UI work. A **design system** is a reusable set of colors, typography, components, and rules that helps a company's products feel consistent. Instead of manually creating every prototype screen, a designer may ask AI to produce a first version of a flow: "Create a three-step account setup process for a college student budgeting app." Then the human designer adjusts hierarchy, accessibility, wording, spacing, and brand fit.

The honest limitation is that an AI-generated prototype can look finished before it is actually good. It may use confusing labels, create inaccessible contrast, ignore edge cases, or produce interactions that are hard to implement. Professional design still requires critique, user testing, and attention to details such as keyboard navigation, screen-reader support, and error states.

9.6 AI in product requirements and team coordination

Product development is full of writing: product briefs, PRDs, user stories, meeting notes, acceptance criteria, bug reports, sprint plans, and launch checklists. AI is now being embedded into collaboration platforms to reduce that paperwork.

Atlassian's 2025 Form 10-K describes Rovo as its advanced AI offering for helping teams locate information, understand it, and take action with specialized agents. It lists Rovo Enterprise Search, Rovo Chat, and Rovo Studio, and says Atlassian has embedded AI capabilities across its platform so customers can get AI-powered productivity enhancements across Atlassian and connected third-party applications. (SEC)

An **AI agent** is software that can use AI to pursue a task, sometimes by taking actions in other tools. In product work, an agent might turn meeting notes into Jira issues, draft a PRD, summarize a design review, or find related customer complaints. But responsible teams keep humans in the approval loop. A PRD written by AI can sound confident while missing the business context, user evidence, or technical constraint that actually matters.

9.7 AI in physical product engineering: generative design and simulation

Physical product development uses a different set of AI tools. **Generative design** is a method where engineers define goals and constraints—such as weight, strength, material, load, connection points, manufacturing method, and forbidden spaces—and software generates possible designs. **Topology optimization** is a related engineering technique that removes or redistributes material while trying to preserve strength or performance.

NASA's Goddard Engineering and Technology Directorate describes its Evolved Structures work as using generative design and digital manufacturing to automate and optimize spacecraft and science-instrument structures. NASA says this improves structural performance by 3x and reduces development time and cost by 10x. Its article explains that a CAD specialist begins by defining mission requirements, connection surfaces, bolts, electronics, optical paths, and areas that must remain open for assembly; then the AI "connects the dots" and can produce complex structure designs in an hour or two. NASA also notes that the algorithms still need a human eye because they can make structures too thin if left unchecked. (Goddard Engineering)

A 2025 open-access engineering study shows the same idea at a smaller scale. Researchers used Autodesk Fusion 360 to generate multiple design variants for a sports go-kart steering wheel under defined loads and manufacturing constraints. Their final 3D-printed ASA polymer prototype achieved a 60% mass reduction while maintaining validated mechanical performance. (ScienceDirect)

In aerospace, GE Aerospace announced in May 2026 that its researchers demonstrated an in-house generative AI app that produced a preliminary design layout for a hypersonic ramjet engine in seconds. GE described this as early design-study work, not a finished certified engine, and said the app allowed engineers to consider multiple flight conditions and scenarios. (GE Aerospace)

The pattern is clear: AI is strong at exploring many options within constraints. Engineers are still responsible for the constraints, validation, testing, safety, manufacturability, and certification.

9.8 AI in digital twins and manufacturing planning

For complex products, development does not stop at the object itself. Teams also need to design the process that will build it. This is where digital twins are important. A digital twin lets teams simulate a product, production line, or factory before committing money to physical changes.

Siemens describes its Xcelerator industrial software portfolio as centered on the digital twin and says a physics-based digital twin combines mechanical, electrical, and software information to help customers make real-world decisions faster and with more confidence. Siemens also says it is applying AI capabilities to design and simulation tasks, such as running many simulations for turbine-blade cooling in the time it previously took to run one. (Siemens Assets)

In October 2025, Siemens and NVIDIA announced a technology stack, still in development, integrating Siemens Xcelerator and NVIDIA Omniverse for advanced factory digital twins. Their press release says the system is intended to bring 3D visualization, simulation, and factory data into one environment and to use AI to simulate hundreds of potential factory layouts. Because this is a vendor partnership announcement, it is evidence of where major suppliers are building capability, not proof that every manufacturer has already achieved those results. (Siemens Press)

9.9 AI in software product development

For digital products, the development stage often means software engineering. AI coding assistants are now one of the clearest examples of AI deployment in product development. These tools help developers autocomplete code, explain unfamiliar codebases, write tests, draft documentation, refactor code, and review pull requests. A **pull request** is a proposed code change submitted for review before being merged into a shared codebase.

Microsoft stated in its FY2025 Q4 earnings call that GitHub Copilot had 20 million users, that GitHub Copilot Enterprise customers increased 75% quarter over quarter, and that 90% of the Fortune 100 used GitHub Copilot. Microsoft also said AI projects on GitHub more than doubled over the previous year and that coding agents such as Claude Code, Codex, Cursor, and GitHub Copilot were generating more pull requests and repositories. (Microsoft)

Independent user sentiment is more mixed. Stack Overflow's 2025 Developer Survey reported that 84% of respondents were using or planning to use AI tools in their development process, with 51% of professional developers using AI tools daily. But the same survey reported that positive sentiment toward AI tools had fallen to about 60%, down from more than 70% in 2023 and 2024. It also found strong resistance to using AI for high-responsibility tasks such as deployment and monitoring or project planning. (Stack Overflow Insights)

Google Cloud's 2025 DORA research surveyed nearly 5,000 technology professionals and found broad AI adoption in software development, with 90% using AI at work and more than 80% saying AI increased productivity. But DORA's central warning is important: AI "amplifies what's already there." Strong teams can get stronger; weak workflows can become more unstable because AI increases the volume of change faster than testing, review, and feedback systems can handle. (Google Cloud)

9.10 How usage differs by business size

Small businesses and startups often use AI because it reduces the cost of first drafts. A founder can summarize customer interviews, create landing-page copy, generate mockups, build a demo, and draft investor or product documents with inexpensive tools. The benefit is speed. The risk is false confidence: a polished prototype can hide weak research, unclear strategy, or insecure code.

Mid-sized businesses usually care about integration. They may already use Jira, Confluence, Figma, GitHub, Adobe, Slack, Microsoft 365, or Google Workspace. For them, AI becomes valuable when it works inside the team's existing information flow: turning support tickets into feature ideas, linking PRDs to engineering tasks, searching internal knowledge, or generating release notes.

Large enterprises care most about control. They need security, privacy, intellectual property protection, audit trails, access permissions, and compliance. That is why company disclosures often emphasize private data controls, enterprise search, model customization, and permission-aware agents. Adobe highlights enterprise-customized Firefly models and brand guardrails; Atlassian describes permission-aware AI across its platform; and Siemens emphasizes digital-thread data and physics-based validation for industrial AI. (Adobe)

9.11 What AI is good at—and what it is not good at

AI is good at accelerating blank-page work. It can produce first drafts, generate alternatives, summarize large text sets, convert rough ideas into mockups, suggest code, and explore engineering design spaces. It is especially useful when the task has many possible answers and humans can judge the output.

AI is weaker when the cost of being wrong is high. It may invent facts, misunderstand users, ignore rare edge cases, leak sensitive information if used carelessly, reproduce bias from training data, or generate code that appears correct but fails under real conditions. In safety-critical fields such as aerospace, medical devices, finance, and automotive systems, AI output must be verified through engineering review, testing, documentation, and regulatory processes.

NIST's AI Risk Management Framework is a helpful reality check. NIST says trustworthy AI should be valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed. Those qualities do not appear automatically just because a tool is impressive; teams have to design review processes around them. (NIST AI Resource Center)

9.12 Hands-on lab: Use AI to design a product concept responsibly

Goal: Create an AI-assisted product concept while keeping humans responsible for judgment.

Time: 60–90 minutes.

Tools: Any general AI chatbot, a document editor, and optionally a design tool such as Figma, Canva, PowerPoint, or paper sketches.

Scenario: Your college wants to reduce long cafeteria lines between 11:30 a.m. and 1:00 p.m. Design a simple product or feature that helps students order, pick up, or choose food faster.

9.12.1 Step 1: Define the problem

Write a two-sentence problem statement without AI first.

Example: "Students lose time waiting in cafeteria lines during peak hours. The college needs a low-cost way to reduce wait times without excluding students who do not want to use an app."

Now ask AI:

Act as a product manager. Given this problem statement, list five user groups, five likely pain points, and five assumptions we would need to test. Keep the language simple.

Review the answer. Circle any assumption that AI invented without evidence.

9.12.2 Step 2: Create research questions

Ask AI:

Create eight interview questions for students and three interview questions for cafeteria staff. The questions should not pressure people toward a specific solution.

Choose the best five student questions and two staff questions. Rewrite them in your own words.

9.12.3 Step 3: Analyze sample feedback

Use this fictional feedback set:

1. "I only have 10 minutes between classes."
2. "The line is long, but I also want to see what food looks fresh."
3. "I tried mobile ordering at another school and my food was cold."
4. "I do not want another app."
5. "The cafeteria staff are working hard; the bottleneck is payment."
6. "I have food allergies and need clear ingredient labels."
7. "I would preorder if pickup times were accurate."
8. "Sometimes the menu online is wrong."

Ask AI:

Group these comments into themes. For each theme, quote the exact feedback numbers that support it. Do not add information that is not in the comments.

Check whether the AI used the evidence correctly. If it creates a theme with no supporting comment, mark it as a hallucination.

9.12.4 Step 4: Generate three concepts

Ask AI:

Based only on the themes above, propose three product concepts. Include one app-based idea, one non-app idea, and one hybrid idea. For each, list benefits, risks, and what we would test first.

Score each concept from 1 to 5 on user value, feasibility, cost, risk, and fairness. The best concept is not always the most exciting one; it is the one that solves the problem under real constraints.

9.12.5 Step 5: Draft a mini PRD

Ask AI:

Draft a one-page product requirements document for the selected concept. Include goal, non-goals, target users, user stories, acceptance criteria, risks, and success metrics.

A **user story** is a simple sentence describing what a user needs, often written as: "As a [type of user], I want [goal], so that [benefit]." **Acceptance criteria** are conditions that must be true before the feature is considered done.

Edit the PRD. Add at least three human corrections.

9.12.6 Step 6: Prototype

Create a rough prototype. This can be a hand sketch, slide, or Figma mockup. Use AI only to suggest layout ideas or wording. Do not let AI be the final judge of usability.

Test your prototype with one classmate. Ask them to complete a task such as "Find today's vegetarian option and choose a pickup time." Record where they hesitate.

9.12.7 Step 7: AI risk check

Use this checklist:

Question	Your answer
What did AI help with?	
What evidence did we actually have?	
What did AI assume?	
What user group might be left out?	
What private data would this product collect?	

Question	Your answer
What could go wrong if the system is inaccurate?	
What must a human approve before launch?	

Your final deliverable is a one-page concept summary, one prototype image or sketch, and the completed risk checklist.

9.13 Key takeaway

AI tools are absolutely being used in product design and development. The best way to understand them is not as replacements for designers, product managers, or engineers, but as accelerators inside the product workflow. They help teams move faster from messy input to structured options. They help create drafts, prototypes, simulations, and code. But the important decisions—what problem matters, what tradeoff is acceptable, what is safe, what is fair, what is legally usable, and what should ship—still belong to accountable human teams.

Appendices

Appendix: AI-Assisted Chapter Development Workflow

This appendix explains how the chapter drafts in this textbook were created with the assistance of AI tools. It is included for transparency, reproducibility, and student learning.

The short version is this: AI tools were used heavily, but not passively. They were used as research assistants, drafting partners, summarizers, structure generators, and revision tools. They were not treated as authorities. Each chapter was reviewed, revised, and aligned by a human instructor before being included in the textbook.

This appendix includes the model and tool setup, the general workflow, the best practices used, and the full markdown prompt used to generate the chapter drafts.

Model and Tool Reference

The chapter drafts were created using **GPT-5.5 Pro** with all tools enabled, especially live web search.

Live search was important because the textbook focuses on how AI tools are being used in business and information technology in **2026**. For this topic, relying only on older training data would be irresponsible. AI tools, business platforms, regulations, public company disclosures, and adoption patterns change quickly. Live search helped locate recent sources, including government reports, company filings, product documentation, regulatory materials, journalism, and empirical research.

The model was used in a tool-enabled environment capable of searching current web sources, reading and comparing sources, drafting long-form instructional text, and assisting with revision.

The model setup can be summarized as follows:

Model: GPT-5.5 Pro **Tool access:** All tools enabled **Especially important tool:** Live web search **Primary task:** Generate introductory college-level chapter drafts on AI tools in business and IT **Drafting style:** Friendly, knowledgeable, evidence-backed, student-facing **Human role:** Instructor, editor, reviewer, fact-checker, curriculum aligner, and final decision-maker

Separate Chat Method

Each chapter draft was generated in a **separate new chat**. The chapters were intentionally isolated from one another during drafting.

This mattered for several reasons. First, separate chats reduced cross-contamination. A chapter about AI in customer service should not automatically inherit examples, assumptions, or source patterns from a chapter about AI in finance or supply chain management.

Second, separate chats made the research question cleaner. Each chapter began from a simple empirical question:

Are AI tools being used in [topic]? How?

For example:

Are AI tools being used in customer service and IT support? How? Are AI tools being used in data mining and analytics? How? Are AI tools being used in digital marketing? How? Are AI tools being used in fraud detection and cybersecurity? How? Are AI tools being used in supply chain management and logistics? How?

That simple question helped keep the chapters grounded. The goal was not to predict a far-off future or repeat marketing language. The goal was to examine current use.

Third, isolated chats made the drafting process easier to audit. Each chapter had its own research path, sources, examples, and hands-on lab. This made it easier to revise or replace one chapter without disrupting the rest of the textbook.

Why Prompt Engineering and Context Engineering Were Used

This textbook used both **prompt engineering** and **context engineering**.

Prompt engineering means writing instructions that guide an AI model toward a desired kind of output. A weak prompt might say: "Write a chapter about AI in business." That would likely produce a generic, shallow, overconfident essay. A stronger prompt defines the audience, level, source standards, structure, tone, task, and evidence expectations.

Context engineering means giving the AI system the right background, constraints, examples, documents, or framing so that its output fits the intended purpose. In this project, the context included the course audience, the VCCS-aligned learning objectives, the textbook's practical orientation, the need for current evidence, and the requirement that every chapter include a major hands-on lab.

Modern knowledge work increasingly involves directing AI tools, checking their output, improving their reasoning environment, and deciding what is good enough to use. This appendix makes that process visible.

Human-in-the-Loop Review

The chapter drafts were produced through a **human-in-the-loop** process. Human-in-the-loop means that a human being remains responsible for reviewing, correcting, approving, or rejecting AI-assisted work. In this textbook, the AI model helped produce drafts, but the human instructor remained responsible for quality.

The human review process focused on several questions:

Did the chapter match the VCCS course outcome? Was the writing appropriate for introductory college students? Were technical terms defined clearly? Did the chapter avoid hype and unsupported predictions? Were sources recent, relevant, and credible? Did the chapter distinguish evidence from vendor marketing? Were real companies and tools discussed carefully? Did the chapter include a useful hands-on lab? Did the lab avoid private data, unsafe instructions, or unrealistic assumptions? Did the chapter explain both benefits and risks? Did the chapter help students build judgment, not just memorize buzzwords?

The AI-generated drafts were therefore treated as strong first drafts, not final authorities. Fluent writing was not considered proof of accuracy. Claims still needed review.

Source Standards Used During Drafting

The chapter prompt instructed the AI model to prefer empirical and verifiable sources over hype.

The general source hierarchy was:

1. **Government data, audits, and standards** Examples include census data, GAO reports, FTC materials, NIST publications, and official regulatory documents.
2. **Public company disclosures** Examples include annual reports, SEC filings, earnings calls, and investor materials. These are not perfect sources, but they are more accountable than ordinary marketing pages because companies make them in formal legal and investor contexts.
3. **Peer-reviewed research and empirical studies** These sources were especially valuable when they measured real-world effects, such as worker productivity, error rates, or adoption patterns.
4. **Independent journalism and industry reporting** Reporting from reputable outlets was useful for public incidents, company deployments, market behavior, and labor impacts.
5. **Vendor documentation** Vendor documentation was useful for explaining what products claim to do, but it was treated carefully.
6. **Vendor case studies and consultancy reports** These were treated with the most caution. They may be useful, but they are often promotional. They were not treated as neutral proof.

This source posture was especially important because AI business writing is often contaminated by marketing language: "transformation," "revolution," "10x," "autonomous enterprise," "hyperpersonalization," and other phrases that sound impressive but do not always explain what is actually happening.

Best Practices Used as AI-Assisted Knowledge Work

The textbook drafting process used several best practices that students can also apply in their own AI-assisted work.

Start with an empirical question

Each chapter began with a concrete question about reality: Are AI tools being used in this area? How? This reduced speculation and kept the model focused on evidence.

Define the audience

The prompt clearly stated that the chapter was for introductory college students, especially community college students with mixed technical backgrounds. This helped control the reading level and required technical terms to be defined.

Require current evidence

Because AI tools change quickly, the prompt required recent sources whenever possible. This is a major difference between AI-assisted knowledge work and ordinary essay generation. For fast-moving topics, current sources matter.

Separate evidence from marketing

The prompt explicitly warned against treating vendor claims, consulting reports, and case studies as neutral evidence. This helped the drafts avoid becoming advertisements for AI products.

Ask for conceptual fundamentals and practice

Each chapter needed to explain concepts and include a hands-on lab. This reflected a core teaching principle: students learn better when they connect ideas to action.

Use AI for drafts, not final truth

The AI model helped create structure, explanations, examples, and lab activities. But the human instructor remained responsible for checking quality, revising language, and deciding what belonged in the final textbook.

Keep humans responsible for judgment

This principle appears throughout the textbook and the drafting process itself. AI can accelerate work, but responsibility remains human. That is true when writing a chapter, reviewing a business forecast, designing a chatbot, or automating an invoice process.

Preserve the prompts

The actual prompt is included below so students can inspect the process. This matters because prompt design is part of the intellectual work. The instructions shaped the output.

General Chapter Drafting Pattern

Although each chapter focused on a different topic, the general process was consistent.

First, a new chat was opened.

Second, the chapter topic was entered as a simple empirical question:

Are AI tools being used in [chapter topic]? How?

Third, the full chapter-writing prompt was supplied.

Fourth, the model conducted live research and drafted the chapter.

Fifth, the draft was reviewed for relevance, accuracy, readability, evidence quality, and alignment with the course outcome.

Sixth, the hands-on lab was checked to make sure it was realistic, safe, and useful for students.

Seventh, the chapter was revised and incorporated into the textbook draft.

The goal was not to hide AI involvement. The goal was to model responsible AI-assisted knowledge work.

Full Markdown Prompt Used for Chapter Drafts

The following markdown prompt was used as the main chapter-generation prompt. For each chapter, the topic was changed to match the relevant VCCS course outcome.

```
# BIT Textbook Chapter Writer

Generate a polished student-facing chapter on how AI Tools are actually deployed in IT
and businesses of any size today (2026), grounded in empirical realities and recent
industry trends rather than consultancy forecasts or vendor marketing. Your voice is
**Friendly Knowledgeable Guide** -- current, evidence-backed, intellectually serious,
and explanatory at an introductory college level. The chapter pairs **conceptual
fundamentals** with **a real hands-on lab**. The required output is one document (~3000
to 3500 words).

The user will typically supply:

- A topic (e.g., "fraud detection in banking," "customer support," etc.

If the topic is not provided, ask one targeted clarifying question before committing
to research. Otherwise, proceed.

## Voice and audience

The chapter is written for **introductory college students** -- mostly community
college, mostly first-year, with mixed technical backgrounds. Calibrate accordingly:

- **Define every technical term on first use.** Assume the student has never seen
technical jargon or acronyms.

- **No heavy math.** Conceptual intuition only. If a formula appears, it's accompanied
by a plain-English translation.

- **Intellectually serious, not dumbed down.** Students notice condescension. Treat
them as capable adults who simply haven't met this material yet.

- **Cite real, current sources.** This skill always researches first; the chapter must
be anchored in named companies, named products, named studies from the last 24 months
wherever possible.

## Source posture: evidence on the ground, not consulting hype

The chapter's job is to describe **what is actually being deployed in 2026**, not what
some consultancy says could happen by 2030. The standing principles are:

- **Prefer disclosure over promotion.** SEC filings (10-K, 10-Q, 8-K), earnings call
transcripts, regulatory submissions, court filings, and government audits beat vendor
blog posts and conference keynotes.

- **Prefer measurement over projection.** A peer-reviewed clinical validation study,
a GAO audit, or a NIST evaluation beats a "TAM" estimate or a "by 2027" forecast.

- **Prefer reporting over thought leadership.** Investigative journalism (Reuters,
Bloomberg, FT, WSJ, ProPublica, The Markup, 404 Media) beats LinkedIn essays and
consultancy "perspectives."

- **Treat the Big Consultancies as interested parties, not neutral observers.** Gartner,
McKinsey, BCG, Bain, Deloitte, PwC, EY, Accenture, IDC, and Forrester all sell AI
consulting services. Their reports — including Gartner's Hype Cycle, Magic Quadrant,
McKinsey's "State of AI," BCG's "AI Maturity Index," and equivalents — are marketing
collateral for their own service lines. They may be cited when **no other source
```

```
documents a particular claim**, and only with explicit framing: "according to McKinsey,
which sells AI consulting services to enterprises, ..." Never treat their adoption
percentages, ROI claims, or market-sizing figures as established fact.

- **Treat vendor case studies as testimonials, not evidence.** A Salesforce customer
story on Salesforce.com is advertising. Use only when the customer has independently
confirmed the deployment (earnings call, SEC filing, press interview, third-party
reporting), and flag inline as "vendor-published."

- **Distrust round numbers.** "40% productivity gain," "50% cost reduction," "10x
faster" are marketing figures unless they appear in a disclosed methodology with
denominators and time windows.

## TASK

Adopting the role of a Subject Matter Expert-Researcher in the AI transformation of
Information Technology & Business (BIT fields), examine, contextualize, and explain at
an introductory college-student level how AI tools are being used for <user provided
topic> in IT and businesses of any size. Format your report as a 3000-3500 well-
structured chapter.
```

Topic Starters Used for Individual Chapters

The general prompt above was reused across chapters. Each chapter began with a topic-specific empirical question. These topic starters followed this pattern:

Are AI tools being used in [topic]? How?

Examples include:

Are AI tools being used in customer service and IT support? How?
Are AI tools being used in data mining and analytics? How?
Are AI tools being used in digital marketing and personalization? How?
Are AI tools being used in fraud detection and cybersecurity? How?
Are AI tools being used in supply chain management and logistics? How?
Are AI tools being used in human resource management? How?
Are AI tools being used in financial forecasting and budgeting? How?
Are AI tools being used in business process automation and robotic process automation? How?
Are AI tools being used in product design and development? How?

These questions were deliberately plain. A more promotional question, such as "How is AI revolutionizing marketing?" would have pushed the model toward hype. A more grounded question, such as "Are AI tools being used in marketing? How?" pushed the model toward evidence.

What I Hope Students Learn From This Appendix

This appendix is not included only to document the textbook's production process. It is also included because the production process is itself an example of the course topic.

Many students will use AI tools in school and at work. The important question is not simply whether AI was used. The important questions are:

Was the use transparent? Was the task appropriate for AI assistance? Was the output checked? Were sources evaluated? Were claims verified? Was private or sensitive information protected? Was a human responsible for the final result? Could someone else inspect the process?

In this textbook, AI was used as a powerful assistant in a structured human-led workflow. That is the same basic pattern students should look for in responsible business and IT uses of AI.

AI tools are not magic knowledge machines. They are systems that can help with searching, summarizing, drafting, organizing, explaining, and revising. Used carelessly, they can produce confident nonsense. Used carefully, they can accelerate serious work.

The difference is not the tool alone. The difference is the process around the tool.