

AI in Operations & Supply Chain Management

Comprehensive Class Notes

From Chatbots to AI Agents — What You Need to Know

BCOR 440: Operations & Supply Chain Management
Gannon University | Spring 2026

The Evolution of AI: Three Waves

Before we dive into tools and techniques, let's zoom out. AI didn't just appear overnight. It evolved through three distinct waves, each fundamentally changing what computers could do for us.

Wave 1: Lookup (Search & Retrieval)

This is the AI most of us grew up with. You typed a question into Google, and it found web pages that probably had the answer. Early chatbots like Siri and Alexa followed the same logic — they matched your words to pre-programmed responses. The AI wasn't *understanding* you. It was pattern-matching.

Think of it this way: Wave 1 AI was like a really fast librarian. It could find the right book, but it couldn't read it for you.

Wave 2: Generation (Creating Content)

Starting around 2022–2023, Large Language Models (LLMs) like ChatGPT, Claude, and Gemini changed the game. Instead of just finding information, AI could now *create* it — writing emails, generating code, summarizing reports, even drafting business plans.

This is where most of us are today. You've probably used an LLM to help with a paper, brainstorm ideas, or explain a concept. The AI isn't just retrieving — it's synthesizing and producing new content based on patterns it learned from massive amounts of text.

The analogy: Wave 2 AI is like a brilliant intern. It can write, analyze, and create — but it can't actually *do* anything in the real world. It can draft an email but can't send it. It can suggest an inventory reorder but can't place the order.

Wave 3: Action (AI Agents)

This is where we are right now, in early 2026. AI isn't just generating content anymore — it's *taking action*. AI agents can browse the web, interact with software, manage files, send messages, and execute multi-step workflows. They don't just tell you what to do. They do it.

The analogy: Wave 3 AI is like a capable assistant who has access to your computer, your tools, and your calendar. You say "Schedule a meeting with the supplier, pull last quarter's inventory data, and draft a summary," and it actually does all three.

⚡ Why This Matters for Your Career

Understanding these waves helps you see where the opportunities are. Most professionals are still in Wave 2 — using AI to generate content. The competitive advantage is in Wave 3: knowing how to **orchestrate AI agents** to handle complex business workflows. That's the skill gap you can fill.

The AI Landscape: Tools You Should Know

The AI ecosystem is evolving fast. Rather than memorizing every tool (they'll change), focus on understanding the categories. Here's a map of what's out there.

Frontier General-Purpose LLMs

These are the “big brains” — models that can handle virtually any text-based task: writing, analysis, coding, research, and reasoning.

Platform	Developer	Key Strengths
Claude	Anthropic	Long-context analysis, nuanced writing, safety-focused design
ChatGPT	OpenAI	Broad general knowledge, large plugin ecosystem, image generation
Gemini	Google	Deep Google integration, multimodal (text + image + video), real-time search

Key insight: These models are more similar than they are different. The real skill isn't picking the “best” one — it's knowing how to use any of them effectively through good prompting and clear thinking.

Specialized AI Tools

While LLMs are generalists, specialized tools are built for specific jobs. They use AI models under the hood but are optimized for particular workflows.

Tool	What It Does	Why It Matters
------	--------------	----------------

ElevenLabs	AI voice generation and cloning	Creates realistic voiceovers for presentations and training materials
Midjourney	AI image generation from text descriptions	Professional-quality visuals without a graphic designer
Cursor / Windsurf	AI-powered code editors	Write and debug software with AI assistance in real-time
NotebookLM	AI research assistant (Google)	Analyzes documents and generates podcast-style audio summaries
Google Vids	AI video creation (Google)	Creates polished videos from text prompts and scripts

How AI Apps Work: The Role of API Keys

Here's something most people don't realize: many of the AI tools you use daily aren't building their own AI. They're *renting* intelligence from frontier models through something called an API (Application Programming Interface).

When an app like Cursor or NotebookLM processes your request, it's often sending your prompt to Claude or GPT behind the scenes, getting a response, and packaging it for you. The API key is like a subscription card that lets the app access the AI's brain.

What This Means for You

Understanding API keys helps you see through the marketing. When a startup says "powered by AI," ask: whose AI? If it's using the same Claude or GPT model you can access directly, the startup's value is in its user interface and workflow — not in its intelligence. Sometimes going directly to the source gives you more power and flexibility.

Key Concepts: Models vs. Agents

This is arguably the most important distinction to understand in AI right now. Let's break it down.

AI Models: The Brain

An AI model (like Claude, GPT, or Gemini) is the underlying intelligence. It's been trained on massive datasets to understand language, reason through problems, and generate responses. But on its own, a model just sits there. It waits for input and produces output. That's it.

A model is like an incredibly smart brain in a jar. It can think, but it can't see, move, or touch anything in the real world.

AI Agents: The Framework

An AI agent takes a model and wraps it in a framework that gives it the ability to *act*. An agent can use tools — browse the web, read files, send emails, interact with software, run code, and make decisions about what to do next.

An agent is like giving that brain a body. Now it can walk around, pick things up, and get things done.

Feature	Description
AI Model	Processes text in → text out. No tools, no actions. Like asking a genius for advice — great answers, but you still have to do everything yourself.
AI Agent	Uses a model as its brain, but adds tools and autonomy. Can search the web, access databases, execute code, send messages, and chain multiple steps together.
Example	Model: "Here's a draft purchase order." Agent: "I've drafted the PO, checked the supplier's inventory system, confirmed pricing, and sent it for approval."

MCP: The Universal Adapter

You might hear the term **MCP (Model Context Protocol)** as AI tools evolve. Here's the simplest way to think about it.

Before MCP, every AI tool needed a custom-built connection to every service it wanted to use. Want AI to work with Slack? Build a Slack connector. Want it to work with Google Drive? Build another connector. This was expensive and fragile.

MCP is like **USB-C for AI**. Instead of building a separate cable for every device, USB-C is one standard that works everywhere. MCP is one standard that lets any AI model talk to any tool or service. This is why AI agents are suddenly getting much more capable — MCP makes it dramatically easier to give them access to the tools they need.

AI That Operates Your Computer

The cutting edge of Wave 3 is AI that can literally use your computer the way you do — clicking buttons, navigating apps, filling out forms, and managing files.

Tool	What It Does
Claude Computer Use / CoWork	Anthropic's system that lets Claude see your screen and interact with desktop applications, automating workflows across any software.
OpenClaw	An open-source framework for AI agents that can operate computer interfaces, designed for developers building automation tools.
Gemini with Project Mariner	Google's approach to browser-based AI agents that can navigate websites and complete online tasks.

Operations Connection

Imagine an AI agent that can log into your ERP system, pull the latest demand forecast, cross-reference it with supplier lead times, identify potential stockouts, draft POs for at-risk items, and flag them for your review — all before your morning coffee. That's not science fiction. That's the trajectory we're on right now.

Understanding Context Windows

Here's something critical that most people miss: AI doesn't "know" everything at once. Every AI model works within a **context window** — a limited amount of text it can process at any given time.

What Is a Context Window?

Think of it as the AI's short-term memory. Everything you put into that window — your prompt, any documents you attach, the conversation history — is all the AI has to work with. It can't see anything outside that window.

Model	Context Window	Roughly Equivalent To
Claude (Opus/Sonnet)	~200,000 tokens	About 500 pages of text
GPT-4o	~128,000 tokens	About 300 pages of text
Gemini 2.0	~2,000,000 tokens	About 5,000 pages of text

Why This Matters

What you put in the window shapes what you get out. If you ask an AI to analyze your supply chain without giving it your actual data, you'll get generic textbook answers. If you feed it your real inventory reports, supplier contracts, and demand forecasts, you'll get specific, actionable recommendations.

This is the direct link to prompt engineering: your job is to fill that context window with the right information so the AI can give you the best possible output.

Pro Tip

When working on complex operations problems, front-load your context. Give the AI the most important information first — the specific data, constraints, and objectives. Don't make it wade through background fluff to find what matters.

AI Hallucinations: When AI Gets Confidently Wrong

This might be the most important section in these notes. AI can and will **make things up**. Not on purpose — it's a mathematical consequence of how these models work. They predict the most likely next word, and sometimes that prediction sounds perfect but is factually wrong.

What Are Hallucinations?

AI hallucinations are outputs that sound confident, fluent, and well-structured — but contain fabricated facts, incorrect data, or invented sources. The dangerous part? They're often indistinguishable from accurate responses unless you know the subject matter.

Real-World Examples

- **Fake citations:** AI models have invented academic papers that don't exist, complete with convincing author names, journal titles, and publication dates.
- **Wrong calculations:** Ask an AI for an EOQ calculation with specific numbers. It might get the formula right but mess up the arithmetic, especially with multiple steps.
- **Outdated information:** An AI might confidently tell you a company's current CEO based on old training data, not realizing leadership changed six months ago.
- **Plausible but false claims:** "Walmart uses a 3-day safety stock for perishables in the Southeast region." Sounds specific and credible. Could be completely made up.

How to Protect Yourself

1. **Always verify critical facts.** If you're making a business decision based on AI output, cross-check the key data points against primary sources.
2. **Ask AI to show its reasoning.** Use chain-of-thought prompting (covered in the next section) to see the logic step by step. Errors are easier to spot when you can see the work.
3. **Be extra cautious with numbers.** AI models are language models, not calculators. They're better at explaining concepts than doing precise math.
4. **Watch for false confidence.** If the AI says something with zero hedging or caveats on a complex topic, that's actually a red flag. Real expertise usually comes with nuance.
5. **Use AI search features.** Models with real-time web search (like Claude, ChatGPT, and Gemini) can ground their responses in current sources. Always prefer search-enabled responses for factual queries.

The Golden Rule

Never present AI-generated content as verified fact without checking it yourself. In a professional setting, you are responsible for the accuracy of your work — whether you wrote it or an AI did. **AI is your first draft, not your final answer.**

Practical Prompt Engineering

Knowing how to communicate with AI effectively is a career skill. It's not about memorizing magic phrases — it's about clear thinking, structure, and context. Here are the techniques that matter most.

1. Prompt Structure: Use Headings for Clarity

When giving AI a complex task, use **# headings** to organize your prompt. This helps the AI parse different parts of your request, just like headings help a reader navigate a document.

Example Prompt:

```
# Role
You are a supply chain analyst at a mid-size retailer.

# Context
We have 3 distribution centers serving 45 retail locations.
Lead times from our primary supplier have increased from 2 to 5 weeks.

# Task
Recommend 3 strategies to reduce stockout risk while minimizing
```

inventory carrying costs. Include pros, cons, and estimated impact.

Why this works: The AI knows exactly what role to play, what information it has to work with, and what output you expect. Compare this to just typing “How do I reduce stockouts?” — the difference in response quality is dramatic.

2. Reverse Prompting

Instead of trying to craft the perfect prompt yourself, ask the AI to help you build it. Tell the AI what you’re trying to accomplish and ask: *“What information would you need from me to give me the best possible answer?”*

The AI will then ask you targeted questions that help you provide exactly the right context. This is especially powerful when you’re tackling a problem you’re not sure how to frame.

3. Meta Prompting

This is one level above reverse prompting. Instead of asking the AI to help with your task directly, you ask it to **design the optimal prompt** for your task. You’re essentially asking the AI: *“If you were an expert prompt engineer, how would you structure a prompt to get the best answer about [your topic]?”*

The AI will produce a carefully structured prompt template that you can then fill in with your specific details. It’s like asking an architect to design the blueprint before you start building.

4. Chain-of-Thought Prompting

For operations problems that involve calculations, multi-step reasoning, or complex trade-offs, ask the AI to **show its reasoning step by step**. This is critical for two reasons.

- **Accuracy:** Models make fewer errors when they work through problems systematically rather than jumping to an answer.
- **Verification:** When you can see every step, you can spot exactly where the logic breaks down — catching hallucinations before they become bad decisions.

Example:

Instead of: “What’s the EOQ for this product?”

Try: “Calculate the EOQ for this product. Show each step of your calculation, explain what each variable represents, and verify your final answer by plugging it back into the total cost formula.”

5. Persona / Role Assignment

Telling the AI *who it is* dramatically changes the quality and relevance of its output. A generic response and an expert's perspective are very different things.

Persona	How It Changes the Response
"You are a lean manufacturing consultant ..."	Focuses on waste elimination, process maps, value streams, and continuous improvement metrics.
"You are a supply chain risk analyst ..."	Emphasizes vulnerability assessment, supplier diversification, contingency planning, and risk scoring.
"You are a hiring manager interviewing for a logistics role ..."	Generates realistic interview questions, evaluates responses, and highlights what employers actually look for.

6. Why Domain Knowledge Is Non-Negotiable

Here's the uncomfortable truth: **GenAI is only as smart as the context you give it.** If you don't know what questions to ask, what data to provide, or how to evaluate the output, AI becomes a fancy autocomplete — producing content that sounds good but may be meaningless or wrong.

This is exactly why you're taking this course. Understanding operations and supply chain concepts gives you the domain knowledge to do three things that AI can't do for you:

1. **Ask the right questions.** Knowing that lead time variability matters more than average lead time for safety stock calculations means you'll prompt the AI with the right variables.
2. **Provide the right context.** Understanding the difference between a push and pull system means you can give the AI the constraints it needs to make relevant recommendations.
3. **Evaluate the output.** If AI suggests a JIT implementation for a company with unreliable suppliers, your domain knowledge tells you that's a terrible idea — even though the AI's reasoning might sound convincing.

Bottom Line

The students who will thrive in their careers aren't the ones who can use AI tools. Everyone will be able to do that. The winners are the ones who **know enough about their domain to use AI tools wisely** — who can prompt with precision, evaluate with expertise, and catch errors before they become costly mistakes.

Career Implications: Your Future with AI

Let's get real about what AI means for your career — regardless of whether you're going into accounting, finance, marketing, management, or MIS.

The Shift: From Executor to Orchestrator

Five years ago, a supply chain analyst spent 3 hours building a demand forecast in Excel. Today, the same analyst spends 20 minutes prompting AI to build it and 40 minutes *interrogating the assumptions*. The value shifted from execution to judgment.

This pattern is playing out across every business function. The repetitive, time-consuming parts of jobs are being automated. What's left — and what's becoming more valuable — is the ability to think critically, make decisions, manage complexity, and orchestrate AI tools to achieve business outcomes.

What Employers Actually Want

Skill Category	What It Looks Like	How This Course Helps
AI Fluency	Comfortable using multiple AI tools; understands models vs. agents; can evaluate AI output	You're learning the landscape, the concepts, and the practical techniques right now
Domain Expertise	Deep knowledge of operations, supply chain, quality, or logistics concepts	Every chapter builds your ability to ask better questions and catch AI errors
Critical Thinking	Can identify flawed reasoning, spot hallucinations, and make sound judgment calls	Practice Operations simulation puts you in decision-making situations where you must evaluate trade-offs
Communication	Can translate between technical AI capabilities and business needs	Pick-a-Hat project requires you to explain operations concepts to different stakeholders
Orchestration	Can design workflows that combine human judgment with AI automation	Course activities mirror the real-world pattern of using AI as a tool while retaining ownership of decisions

The Architect Mindset

The most valuable professionals in the next decade won't be the ones who can use one AI tool really well. They'll be the ones who can **architect solutions** — combining the right AI tools, the right data, and the right human judgment to solve complex business problems.

That means knowing when to use AI and when not to. Knowing which tool fits which problem. Knowing how to validate outputs. And knowing how to communicate AI-generated insights in ways that drive action.

You're not just learning operations. You're learning how to think about business problems in a way that makes you *irreplaceable* — because you bring the judgment, context, and strategic thinking that AI still can't replicate.

Your Action Items

1. Start using at least two different LLMs this week (e.g., Claude and Gemini). Notice how they differ.
2. Practice structured prompting with # headings on your next assignment.
3. Try reverse prompting: Ask the AI what it needs from you before giving it a task.
4. When AI gives you an answer, always ask: "How do I know this is right?"
5. Build your LinkedIn profile with the AI tools and certifications from this course.

Quick Reference: AI Terminology

Keep this handy as you encounter AI concepts throughout the course and in your career.

Term	Plain English Definition	Operations Example
LLM (Large Language Model)	An AI trained on massive text data that can understand and generate human language	Using Claude to draft supplier evaluation criteria
Context Window	The amount of text an AI can "see" and work with at one time	Uploading a 50-page inventory report for analysis
Hallucination	When AI generates confident but factually incorrect output	AI inventing a supplier name or fabricating lead time data
Prompt Engineering	The skill of writing effective instructions for AI	Structuring a request with role, context, and task sections

API (Application Programming Interface)	A connection that lets one software talk to another	An inventory app using GPT to generate reorder recommendations
MCP (Model Context Protocol)	A universal standard for connecting AI to external tools	AI agent accessing your ERP, email, and calendar through one protocol
AI Agent	An AI system that can use tools and take actions autonomously	Software that monitors inventory levels and places orders automatically
Token	The basic unit AI uses to process text (roughly 3/4 of a word)	Determines how much data you can feed into a single AI interaction
Fine-tuning	Customizing a general AI model for specific tasks or industries	Training a model on your company's historical demand data
RAG (Retrieval-Augmented Generation)	Giving AI access to specific documents to ground its responses in real data	Connecting AI to your supply chain database for accurate answers

*Remember: AI is the most powerful tool you'll use in your career.
But a tool is only as good as the person using it.
That's why you're here.*