BCOR 440 / IE 425 • Operations & Supply Chain Management

Student Field Guide

The AI Orchestrator's Playbook

Agents, Tools & Real-World Operations

Claude Code • Claude Cowork • OpenClaw

Part 1	Part 2	Part 3
Understanding AI Agents	Your Three AI Tools	Operations & Data Analytics

PART 1: Understanding AI Agents — Brain vs. Body

What Is an AI Agent? (And Why It's NOT an AI Model)

Before we touch any tool, we need to get one thing crystal clear — because almost everyone gets this wrong.

Most people think "AI agent" is just a fancy name for ChatGPT or Claude. It's not. Not even close. Confusing an AI model with an AI agent is like confusing a car engine with a car. One thinks. The other acts.

AI Model (The Brain)

This is the Large Language Model (LLM) — the intelligence engine. It reads your prompt, reasons through it, and generates a response in text. It cannot touch your computer, browse the web, or take any action in the real world. It just thinks and writes. Examples: Claude 3.5 Sonnet, GPT-4o, Gemini 1.5 Pro.

AI Agent (Brain + Body)

This is a software framework that wraps around an AI model and gives it hands, eyes, and memory. The agent can browse websites, read and write files, run code, send emails, and take real actions — all by directing the AI model to decide what to do next.

The Hiring Analogy That Makes This Click

> **Professor's Analogy**
> "Hiring an AI agent is like hiring a human assistant. The agent framework is the assistant's body and workflow — but you still have to give them a computer, office supplies, and access to a library so they can actually think and do the work. The agent is the assistant. The AI model is their brain. You need both."

Let's make this even more concrete with a supply chain example:

Scenario	Chatbot (Brain Only)	AI Agent (Brain + Body)
Track supplier delays	"Here are 5 strategies to track supplier delays..."	Logs into your supplier portal, pulls delay data, flags orders at risk, emails your procurement manager.
Analyze inventory data	"Here's how you calculate reorder points..."	Opens your Excel file, calculates reorder points for every SKU, highlights stockout risks in red, saves the updated file.
Competitive pricing	"Amazon uses dynamic pricing algorithms..."	Scrapes competitor prices from 10 websites, compares to your price list, generates a recommendation report.

Weekly demand forecast	"Demand forecasting uses methods like ARIMA..."	Pulls your sales history, runs the forecast model, creates a chart, drops it in your shared folder every Monday at 8 AM.

The Observe → Think → Act Loop

Every AI agent — no matter which tool you use — operates on the same fundamental cycle. Understanding this loop is what separates people who use AI from people who orchestrate it.

1. OBSERVE	2. THINK	3. ACT	4. REPEAT
The agent looks at its environment — your files, a webpage, an API, your instruction.	The AI model (the brain) decides what action to take next to get closer to the goal.	The agent executes — it clicks, writes code, calls an API, moves a file.	It checks the result and loops back to Observe until the goal is achieved.

> **Why This Matters for You**
> This loop is why AI agents are so powerful — and why they require an orchestrator. Without someone designing the goal, the guardrails, and the success criteria, the loop can go wrong. That's your job.

The Three Layers of an Agent System

When you set up any AI agent tool, you're actually building three layers at the same time. Most beginners only think about Layer 1 and wonder why things go wrong.

Layer	What It Is	Your Responsibility
Layer 1: The Brain (LLM)	The AI model doing the reasoning — Claude, GPT-4o, Gemini, etc.	Choose the right model for the task (speed vs. accuracy vs. cost)
Layer 2: The Body (Framework)	The agent software: Claude Code, Claude Cowork, or OpenClaw	Configure it, give it the right tools and permissions
Layer 3: The Guardrails (YOU)	The rules, approvals, and checkpoints that keep it from going off the rails	Define what the agent CAN do, what it CANNOT do, and who reviews what

Chatbot vs. Agent: A Side-by-Side Reality Check

Chatbot (e.g., Standard Claude, ChatGPT)	AI Agent (Claude Code, Cowork, OpenClaw)
Responds only within the chat window	Can interact with your files, browser, and external apps

No memory between conversations (by default)	Persistent memory — it remembers you across sessions
You copy-paste its output into your work	It does the work directly in your environment
Great for brainstorming and drafting	Great for automating, executing, and building
Forgets everything when you close the tab	Runs in the background, can work while you sleep
Free or low-cost entry point	Requires API keys, setup, and thoughtful configuration

The “Bring Your Own Brain” Concept

Here’s one of the most important insights in the whole AI agent ecosystem: the agent framework and the AI model are separate things. You can mix and match them.

Tool (The Body)	Which Brains It Accepts	Best For
Claude Code	Anthropic models only (Claude 3.5 Sonnet, Claude 3 Opus)	Coding tasks, file management, technical projects
Claude Cowork	Anthropic models only	Documents, research, collaborative projects
OpenClaw	Any model: Claude, GPT-4o, Gemini, local models	Custom automation, web scraping, mixed workflows

> **Pro Tip: Model Economics**
> OpenClaw’s “Bring Your Own Brain” flexibility is a huge advantage. Need heavy reasoning? Use Claude Opus. Doing a simple, repetitive scraping task? Use a cheap, fast model to save money. You’re now making architectural decisions — that’s exactly what an AI orchestrator does.

PART 2: Your Three AI Tools — A Beginner's Field Guide

Meet Your AI Coworkers

You've got three tools in your toolkit this semester. Each one has a distinct personality, a best use case, and a learning curve. Here's how to think about each one — and how to actually start using them.

Claude Code

The Developer's Power Tool — Your terminal-based AI teammate

- Built by Anthropic — lives inside your computer's terminal (command line)
- Reads your entire project folder natively — no copy-pasting code back and forth
- Can write, run, test, and fix code autonomously
- Commits changes to files on your behalf (with your approval)
- Best for: data analysis scripts, Excel automation, building tools for your projects
- Brain: Anthropic models only (Claude 3.5 Sonnet recommended)

Claude Cowork

The Project Manager's Dream — Visual AI workspace for knowledge work

- Beautiful visual interface — no terminal required
- Reads, synthesizes, and organizes documents (PDFs, Word files, research papers)
- Creates task boards, drafts content, and manages projects collaboratively
- Best for: research synthesis, group project planning, report drafting, literature reviews
- Brain: Anthropic models only
- Think of it as: Google Docs + Trello + a brilliant research assistant, all in one

OpenClaw

The Open-Source Multi-Tool — Scrappy, flexible, and infinitely customizable

- Open-source (built by the community, runs on your own computer)
- Connect via WhatsApp, Telegram, Discord, iMessage, or Slack — message it like a coworker
- Persistent 24/7 memory — knows your preferences, your context, your world
- Browses the web, fills forms, scrapes data, controls your computer
- Best for: web automation, data collection, connecting multiple apps, custom workflows
- Brain: YOUR CHOICE — Anthropic, OpenAI, Google, or local models via API key
- Formerly known as Clawdbot and Moltbot (same product, evolved name)

Getting Started: Step-by-Step for Absolute Beginners

Claude.ai — Start Here First (Week 1)

Before you touch any agent tool, spend your first week just using Claude.ai at claude.ai. Think of this as your warm-up lap. You're learning to communicate with AI before you hand it real power.

1. Go to claude.ai and create an account (free tier is fine to start)
2. Try the Chatbot vs. Agent comparison: ask Claude "how do I organize a messy supplier database?" and compare that to what an agent could actually DO with that database
3. Practice prompt engineering: take one bad prompt, rewrite it 3 times, compare the outputs
4. Deliberately break it: ask something it can't know (today's stock prices). Notice the limitation. This is a feature, not a bug — knowing AI limits is a core orchestrator skill

> **Week 1 Checkpoint**
> Goal for Week 1: You understand what a good prompt looks like, you've experienced AI failure firsthand, and you can articulate the difference between a chatbot and an agent in your own words.

Claude Cowork — Your Visual Agent (Week 2–3)

Claude Cowork is the friendliest entry point into actual agent work because there's no terminal involved. It works like software you already know.

5. Download the Claude desktop app from claude.ai — Cowork is accessible through it
6. Create a new project space and drag in a few documents (try your course textbook chapters or case study PDFs)
7. Give it a goal: "Read these three PDFs and create an outline comparing how each company manages inventory differently"
8. Watch it work — don't interrupt. Observe the Observe-Think-Act loop in action
9. Critique it: what did it get right? What did it miss? Tell it to revise. This is iteration, and it's how real orchestrators work

> **The Key Moment**
> You just delegated a task to an AI agent instead of just asking it a question. That's the mindset shift that matters. You're a manager now, not just a user.

Claude Code — Terminal-Based Coding Agent (Week 4–5)

Claude Code has the steepest learning curve of the three because it requires you to be comfortable with the terminal. But business students who can use it become 3x more valuable than those who can't.

10. Step 0 — Get terminal-comfortable: Spend 30 minutes on a YouTube tutorial for Mac Terminal or Windows Command Prompt. You need to know: how to navigate folders (cd), list files (ls or dir), and run a command. That's it.
11. Step 1 — Install Node.js: Go to nodejs.org, download, and install. This is a free software environment that Claude Code needs to run.
12. Step 2 — Install Claude Code: Open your terminal and type: npm install -g @anthropic-ai/claude-code
13. Step 3 — Authenticate: Type claude to start. You'll log in with your Anthropic account.
14. Step 4 — First real task: Navigate to a folder that has a spreadsheet or CSV file. Type: claude "Analyze this data file and tell me the top 3 insights an operations manager should care about"
15. Step 5 — Watch and approve: Claude Code will propose actions before executing them. Approve each one. Never auto-approve in your first week.

> ⚠ **Critical Safety Rule**
> The golden rule for beginners: Claude Code will ask for permission before taking actions. ALWAYS read what it's proposing before clicking approve. You're the manager. The agent is the intern. Review everything.

OpenClaw — Your Personal AI Agent (Week 5–6)

OpenClaw is the most powerful and most flexible of the three tools — but it requires the most setup. Think of it as building your own custom AI assistant that lives on your computer 24/7 and that you can message from WhatsApp or Telegram.

16. Install via one line in your terminal: curl -fsSL https://openclaw.ai/install.sh | bash (Mac/Linux) or follow the Windows instructions at openclaw.ai
17. Run the onboarding: openclaw onboard — it will walk you through naming your agent and connecting an AI brain via API key
18. Get an API key: You'll need an Anthropic API key (from console.anthropic.com) or an OpenAI key (from platform.openai.com). Note: API keys have small costs per use — budget $5–10 for the semester to experiment
19. Connect a messaging app: Link OpenClaw to your Telegram or WhatsApp so you can message your agent like a coworker
20. First task: Tell your agent to monitor a specific website and send you a message if anything changes. This is your first real automation workflow

> **OpenClaw Safety First**
> Always run OpenClaw in dry-run mode for new tasks first. This shows you what it PLANS to do without actually doing it. Review the plan. Then approve. Supply chain analogy: this is your quality check before the product ships.

When to Use Which Tool	Best Fit
You need to analyze a CSV or Excel file of operations data	Claude Code

You’re synthesizing 10 research papers for a project	Claude Cowork
You want to scrape pricing data from competitor websites	OpenClaw
You’re drafting a supply chain risk report collaboratively	Claude Cowork
You need a Python script to calculate EOQ for 500 SKUs	Claude Code
You want to get daily briefings via WhatsApp about market changes	OpenClaw
You’re organizing a group project with multiple deliverables	Claude Cowork
You want to automate weekly inventory reports	Claude Code or OpenClaw

PART 3: AI Orchestration in Operations & Supply Chain

Why Operations Students Have a Superpower Here

Here's a secret most tech people don't know: operations thinking is the best preparation for AI orchestration. Why? Because you already know how to think about processes, bottlenecks, capacity, and failure modes.

Every concept you've learned this semester maps directly to AI agent design. You just need to see the connection.

Operations Concept	AI Orchestration Equivalent	Why It Transfers
Process mapping	Agent workflow design	Before you can automate anything, you must deeply understand the process — same skill
Bottleneck analysis	Finding where human judgment is essential	Where does the agent slow down? Where does it need a human in the loop?
Quality control	Output validation and error checking	How do you know if the agent's work is accurate? You need inspection points
Safety stock	Agent guardrails and limits	Buffer against uncertainty — define what the agent cannot do without approval
Continuous improvement	Agent iteration and refinement	Measure output quality, identify failures, refine the prompt or workflow
Supplier risk management	Dependency mapping in multi-agent systems	What if one part of the system fails? How does the whole workflow respond?

AI Agents in Action: The Four Course Zones

You know the four zones that organize this course. Here's exactly how AI agents plug into each one — with real tasks you could automate today.

Zone 1: How We Make Things — Processes, Projects, Capacity

What Agents Can Automate

- Process flow diagrams — input a text description, Claude Code generates a visual flowchart

Where Humans Must Stay in the Loop

- Strategic capacity decisions (build a new facility vs. outsource)
- Trade-off calls that involve workforce, community, or ethical impact

- Capacity utilization calculations across multiple production lines
- Project schedule generation from scope documents (Gantt chart via Excel)
- Bottleneck identification from production data files

- Final project approval and stakeholder sign-off

Zone 2: Where Stuff Comes From — Supply Chain, Sourcing, Lean

What Agents Can Automate

- Supplier performance monitoring — OpenClaw checks supplier portals daily and flags late deliveries
- Lean waste identification — analyzing process logs to spot transportation, motion, or wait waste
- Supplier research and benchmarking from public data sources
- News monitoring for supply chain disruption signals (raw material price spikes, port strikes)

Where Humans Must Stay in the Loop

- Supplier selection and contract negotiation (relationships and ethics matter)
- Geopolitical and reputational risk assessment
- ESG and sustainability compliance judgment calls

Zone 3: How We Plan — Forecasting, Inventory, MRP

What Agents Can Automate

- Demand forecasting — Claude Code runs time-series models on your sales history every week
- EOQ and ROP calculations across entire product catalogs (not just one SKU)
- ABC analysis — classify thousands of SKUs by velocity and value automatically
- Automated reorder triggers when inventory hits safety stock thresholds

Where Humans Must Stay in the Loop

- Judgment calls when the forecast doesn't match reality (new product launch, crisis)
- Final purchase order approval above a dollar threshold
- Strategic inventory investment decisions (buy vs. make vs. lease)

Zone 4: How We Optimize — Quality, Scheduling, Constraints

What Agents Can Automate

- Statistical process control charts — Claude Code generates control charts from quality data files automatically
- Defect pattern analysis — spotting which machines or shifts are causing quality issues
- Schedule optimization across multiple constraints (machines, workers, due dates)

Where Humans Must Stay in the Loop

- Final root cause determination (AI drafts the analysis; humans validate it)
- Customer communication when quality failures occur
- Recall decisions and regulatory reporting

• Root cause analysis first drafts from incident reports and defect logs	

PART 4: AI Agents for Data Analytics & Data Management

You're Not a Data Scientist — But You Can Think Like One

Here's the truth: most business decisions live or die on data. And most data analysis at real companies is done by people who are NOT data scientists — they're operations managers, supply chain analysts, and business graduates who learned to ask the right questions of the right tools.

AI agents just collapsed the gap between "I have a spreadsheet" and "I have an insight" from hours to minutes. Here's how to close that gap.

The Data Analytics Workflow — AI Agent Style

Stage	Traditional Approach	With AI Agent
1. Data Collection	Manually export reports, copy-paste from multiple systems, 2–3 hours	OpenClaw or Claude Code scrapes, pulls from APIs, and merges data automatically
2. Data Cleaning	Finding and fixing errors, formatting issues, duplicates — the least fun part	Claude Code identifies and fixes messy data, flags anomalies for human review
3. Exploratory Analysis	Pivot tables, basic charts, looking for patterns manually	Claude Code generates summary stats, correlation analysis, and charts in seconds
4. Interpretation	You figure out what the data means	Agent drafts hypotheses; YOU validate them with domain expertise — non-negotiable
5. Reporting	Format charts, write narrative, build slide deck — hours of work	Claude Cowork synthesizes findings into a structured report draft
6. Distribution	Email the report to stakeholders manually	OpenClaw sends it automatically to the right people at the right time

Practical Prompt Templates for Operations Data Analysis

These are ready-to-use prompts you can adapt for Claude Code. Notice the pattern: every good prompt includes CONTEXT, a specific GOAL, the DATA FORMAT, and what you want the OUTPUT to look like.

> **Prompt Template 1: ABC Analysis / Stockout Identification**
>
> Context: I am a supply chain analyst at a mid-size retail company. Goal: Analyze this inventory CSV file and identify the top 20% of SKUs that account for 80% of our stockout events. Data: The file has columns for SKU ID, product name, category, stock level, reorder point, and

stockout date. Output: Create a new Excel file with the flagged SKUs highlighted in red, sorted by stockout frequency, with a summary table at the top.

Prompt Template 2: Forecast Accuracy Review

Context: I am preparing a weekly operations review. Goal: Compare our actual vs. forecasted demand for the past 12 weeks using this Excel file. Calculate Mean Absolute Percentage Error (MAPE) for each product category. Data: Columns include week number, product category, forecasted demand, and actual demand. Output: Generate a bar chart showing forecast accuracy by category, and flag any category with MAPE above 15% as needing attention. Save as a PDF report.

Prompt Template 3: Quality Defect Pattern Analysis

Context: I am investigating a quality control issue at our production facility. Goal: Analyze this CSV of defect records from the past 6 months and find patterns — which shift, which machine, which day of week, and which operator ID is associated with the highest defect rates. Data: Columns include date, shift (AM/PM), machine ID, operator ID, defect type, and defect count. Output: Write a plain-language summary of the top 3 findings, and create a pivot table showing defect rates by all four variables.

Prompt Template 4: Supplier Scorecard

Context: I am building a supplier scorecard. Goal: Using these three CSV files (on-time delivery rates, quality defect rates, and invoice accuracy rates for our top 50 suppliers), calculate a composite performance score for each supplier. Weight: delivery 40%, quality 40%, invoice accuracy 20%. Output: Create a ranked list of suppliers with their composite score, flag anyone below 70 as at-risk in red, and generate a summary for our quarterly business review.

Data Management: Teaching Agents to Handle Your Data Responsibly

This is where the orchestrator's judgment really matters. AI agents can do a lot with your data — but data mismanagement is a real risk. Your job is to design guardrails before any data agent runs.

Risk	What Can Go Wrong	Your Guardrail Design
Data quality	Agent produces analysis on dirty data — garbage in, garbage out	Always tell the agent to validate data first and report anomalies before analyzing
Data privacy	Agent uploads sensitive customer or employee data to an external service	Never give an agent access to PII without reviewing data sharing policies first
Overwrite errors	Agent modifies the original file instead of creating a copy	Always instruct the agent: "Work on a copy, never the original file"
Hallucinated insights	Agent confidently reports a trend that doesn't actually exist in the data	Cross-check key findings manually; use spot-checks on 10% of the output

Scope creep	Agent modifies more files than intended while “cleaning things up”	Specify exact file names and folder paths; use read-only permissions where possible

The Five Golden Rules of Agent Whispering

Managing an AI agent is more like managing a smart but literal-minded intern than programming a computer. These five rules will save you from 90% of beginner mistakes.

1	**Give context, not just instructions.** Don’t say “analyze this data.” Say “I’m an operations analyst preparing a quarterly review for our CFO. Analyze this inventory file and highlight the three most important trends.” Context changes everything.
2	**Define done.** Agents can loop forever if they don’t know when to stop. Always tell them what success looks like: “Stop when you have identified 10 at-risk suppliers” or “Save the output and then stop.”
3	**Human in the loop for every critical decision.** The agent can draft the purchase order. You click send. The agent can flag the defective batch. You decide whether to halt production. Never let an agent make irreversible decisions without your approval.
4	**Embrace iteration.** Your agent will fail. It will misinterpret. It will get partway and get stuck. Don’t start over — just correct it: “That missed the point, the goal is X not Y, try again.” That’s management.
5	**You take the blame if it messes up.** You are the manager. The agent is the intern. If an AI agent sends a wrong email or deletes the wrong file, the professional accountability is yours. Own it. Design your guardrails accordingly.

PART 5: Becoming an AI Orchestrator — Your Career Roadmap

The ORCA Framework: Your Orchestrator's Decision Tool

Every time you design or manage an AI agent workflow, run it through ORCA. This four-step mental model will keep you from making the mistakes that get people fired.

	Step	The Question	Operations Example
O	**Observe the Task**	What is the input? What is the expected output? What does a good result look like?	Input: daily inventory CSV. Output: flagged SKUs below safety stock. Good result: zero false positives.
R	**Route It**	Automate (rule-based), Augment (human + AI), or Delegate to an Agent (full goal)?	Reorder triggers → Automate. Supplier negotiation strategy → Augment. Weekly demand report → Agent.
C	**Configure Guardrails**	What can the agent do autonomously? What requires your approval? What should never happen?	Can: flag orders. Cannot: place orders over $10K without approval. Never: contact suppliers directly.
A	**Assess the Output**	Did it work? Are the outputs accurate? Are humans overriding it too often (red flag)?	Check: did the flagged SKUs match the ones the buyer already knew about? Any surprises to investigate?

What Does This Look Like in a Job Interview?

This is where everything comes together. The students who land the best offers aren't just saying "I know how to use AI." Everyone says that. Here's what the orchestrators say instead.

✗ The Tool User's Answer

"Yes, I use ChatGPT and Claude to help me write emails, summarize documents, and answer questions. I've been using them for about a year."

☑ The Orchestrator's Answer

"In my operations course, I designed a workflow where an AI agent monitored our simulated supplier database, flagged at-risk orders using criteria I defined, and generated a weekly summary report without my involvement. I built the guardrails, tested for failures, and iterated on the design. I can walk you through the architecture if you'd like."

Your 4-Month Skill Progression Roadmap

Month	Focus	What You Can Say in an Interview
Month 1	Claude.ai power user. Practice prompt engineering daily. Deliberately break it and understand why.	"I understand how LLMs reason, where they fail, and how to write prompts that get reliable outputs."
Month 2	Claude Cowork. Build one real project from scratch using the agent. Synthesize course materials.	"I've used AI agents to manage and synthesize large volumes of documents, producing actionable summaries."
Month 3	Claude Code. Take a real dataset from class. Direct the agent to clean, analyze, and visualize it.	"I directed an AI agent to analyze an inventory dataset and identify reorder risk — without writing any code myself."
Month 4	OpenClaw or integration project. Automate one real recurring task in your life or a student org.	"I built and deployed a personal AI automation that saves me X hours per week. Here's how I designed the guardrails."

Your AI Orchestrator Portfolio Checklist

By the end of this course, you should be able to point to tangible evidence of orchestration — not just tool use. These are the portfolio pieces that matter.

Portfolio Piece #1: The Bot Brief

1 Bot Brief: A one-page document describing an AI agent you designed for a real operations problem. Includes: what it does, the data it needs, what it decides autonomously, what triggers human review, and what a bad output looks like.

Portfolio Piece #2: AI-Assisted Data Analysis

1 Data Analysis Project: A real dataset analyzed with Claude Code. The output includes cleaned data, at least one visualization, and a plain-language summary of the top 3 insights for an operations manager. You didn't write the code — you directed the agent that wrote it.

Portfolio Piece #3: A Live Automation

1 Automated Workflow: A recurring task that you've set up an agent to handle. Could be weekly reporting, supplier monitoring, news alerts, or data collection. Document the ORCA framework you used to design it.

Portfolio Piece #4: The Failure Analysis

1 Failure Post-Mortem: A short write-up of a time your agent got it wrong. What happened? Why? What guardrail would prevent it next time? This shows more maturity than listing only your successes — and interviewers love it.

The Bottom Line

A manager who uses AI is like a driver. An orchestrator who designs AI systems is like someone who builds the road. Which one do you want to be? The answer to that question is what this guide is preparing you for. Not just using tools — designing systems. Not just prompting — orchestrating. Not just a user of AI — a manager of AI workforces. Your operations background isn't a detour from this future. It's the fast lane.