

COMMUNITECH · JUNE 18

Build with AI



Tech Worker Edition

Thursday, June 18 · 5:30–8:30 PM EDT

Presented by Momentum Product Co. in
partnership with Communitech



DO THIS NOW

Get Your Assets & Open Your AI Tool

Workshop Assets

All participant guides, prompts, and sample data are at:

momentum-workshops.vercel.app/build-with-ai/

Open this now — you'll need it for every module.

What's inside:

- Participant guides for all 3 modules
- Sample data for Module 1
- Prompt templates
- Set-up instructions

Your AI Tool

Pick one — both work for tonight:

Claude.ai

Browser-based. Go to claude.ai — no install, works now.

You will likely run out of context if you on are the free version.

Cowork (Recommended)

Claude's desktop app. Best for Module 1 — handles file uploads and connectors natively.

Momentum Product Co

We help B2B SaaS companies navigate product and market inflection points — when strategy, execution, and competitive pressure are out of sync — by stepping in as **senior advisors or fractional leaders**.

Our focus is not theory or experimentation theatre. It's deciding what to build, how to accelerate velocity, and how to move faster than competitors — often by unlocking [underutilized AI capabilities](#) inside the organization.

AI-Powered Discovery

[KNOW WHAT TO BUILD BEFORE YOU BUILD IT](#)

10x faster customer insight gathering and market analysis.

Signal over gut-feel: AI-augmented prioritization.

Operating Model & Enablement

[MAKE AI THE DEFAULT, NOT THE EXCEPTION](#)

AI-native workflows across Product, Design, and Engineering.

From occasional experiment to daily muscle memory.

Strategy & Experimentation

[STOP DEBATING. START TESTING.](#)

Cut through the noise to find where AI creates real leverage.

Validate bets in weeks, not quarters. Vibe-code for validation.

Competitive Intelligence

[SEE WHAT YOUR COMPETITORS ALREADY SEE.](#)

Automated monitoring across every channel that matters.

Synthesized signals that shift your roadmap.

👤 YOUR INSTRUCTOR

Heather Gawel

Founder & Product Practice Lead · Fractional Product Leadership

AI-First Product Leadership

Former VP of Product & Design and Innovation Lab Leader. Two decades guiding teams through 0→1, post-PMF scale, and \$100M+ portfolios — at the intersection of AI, product strategy, and execution.

Momentum combines senior product leadership with hands-on AI enablement to transform how your product and design teams work — not what they ship, but how they discover, prioritize, and validate. We get your team to the first 20%. After that, they run with it.



Course Support

Technical AI Lead, Momentum Product Co.



Babajide Okusanya

Product Leader, Software Engineer & AI Consultant / Kitchener

Helps enterprise and startup engineering teams actually adopt AI — not just talk about it. Trained **2,000+ developers** on AI-assisted workflows. Builds daily with Claude Code, Lovable, and N8N. Previously co-founded a B2B supply chain startup. Product work across SaaS, immigration tech, consulting, and nonprofit.

What he nerds out about:

- Spec-driven AI development (not vibe coding)
- Building disposable tools fast — purpose-built AI apps for specific workflows
- Multi-agent orchestration and how teams can actually ship with AI

BEFORE WE DIVE IN

How to Get the Most Out of Tonight

Use your real work (if you can)

Your own transcripts, tickets, and competitors teach better lessons than generic examples.

Messy is expected. Fast is the point.

Bad prompts and wrong outputs are part of the process. Fix on the fly and keep moving.

Hands-on is non-negotiable.

The value is in building alongside the demos — not just watching. Open your laptop.

Some of this will feel basic. Stay with us.

Mixed skill levels, same material. If it feels easy, ask us how to level it up.

Ask when you're stuck.

We'll be walking around. Wave us over — your question is probably everyone's question.

You'll leave with three working artifacts.

A synthesis workflow, a competitive intel agent, and a prototype. All built by you, tonight



5:50 PM

MODULE 1 OF 3

Build with AI: Module 1

Discovery & Customer Insight Synthesis



"Your task: synthesize these sources into a pricing & packaging recommendation — should we change pricing, and if so, how?"

This is about signal, not just synthesis:

Validate your assumptions instead of confirming them

Don't guess - draw on actual data and let the evidence change your mind.

Surface buried signal for the people downstream of your decisions

Build influence through strategic presentation of data.

Apply human judgment

Claude counts frequency; you decide what matters. High frequency does not necessarily mean high impact.

The Problem This Workflow Solves

Pricing gets set by gut

Then it never gets revisited
— even when the signal is
right there.

The data already exists

Customer interviews, tickets,
NPS — it's all sitting unused.

AI collapses the timeline

Days of synthesis → minutes. You pick the question. You own the
call.

This exercise changes that. Take a realistic set of customer data and use Claude to synthesize it into a structured framework — ready for a product strategy conversation — in **under 40 minutes**.



Why Cowork for This Module

Module 1 runs entirely inside **Cowork** — Claude's desktop app. No terminal. No copy-paste marathon. Files load by upload or connector (Gong, Zoom, Google Drive), and it's the exact setup you'll reuse on your own data the day after the workshop.

1

No Terminal Required

- Works entirely in the desktop app interface
- Upload files directly or connect via Gong / Zoom / Drive connectors
- No copy-paste — point Cowork at a folder and it reads everything
- Ideal for live workshop settings and day-to-day reuse

2

Connectors Load Data Directly

- Granola, Gong, Zoom, Google Drive — no exports needed
- Pull discovery data together without leaving the app
- Connectors make the workflow repeatable on new data every sprint
- Same setup you'll operationalize on your own customer data

3

Files Stay Local

- Sensitive transcripts never leave your machine
- Upload CSVs, paste raw text, or attach exports directly
- Works with any format: Granola notes, Zendesk CSVs, NPS exports
- Clean, visible context — you see exactly what Claude sees

 Ready to script it? Graduate to the CLI to automate and version-control the workflow later.

What You'll Need & The Sample Data

All files are in the zip — download it from the Course Files before we start.

Cowork (Recommended)

If you have Cowork access tonight.

1. Download Module 1 · Cowork starter (.zip)
2. Unzip it on your Desktop
3. In Cowork, point at the unzipped folder as your workspace
4. Open 00_START_HERE.md and follow it

Claude Projects (Fallback)

If you don't have Cowork tonight.

1. Download Module 1 · Claude Projects (.zip)
2. Unzip it on your Desktop
3. In Claude, create a new Project
4. Drag everything inside Upload_to_Files/ into the Project's Files panel
5. Paste the contents of PASTE_into_Instructions.txt into the Project's system instructions
6. Start prompting



MODULE 1 OF 3

Module 1: Essential Files

This module focuses on pricing analysis. The zip file contains all the necessary data and tools to conduct your exercise tonight.

- **Customer Discovery Interview Transcripts (Feb–Mar 2026)**

Ten qualitative interviews providing raw customer insights into needs, challenges, and value perception.

- **NPS Scores · Q1 2026 (CSV)**

Quantitative Net Promoter Score data to measure overall customer satisfaction and loyalty.

- **Support Tickets · Q1 2026 (CSV)**

A dataset of customer support interactions, highlighting common issues, feature requests, and user feedback.

- **Launchpad Pricing & Win/Loss Supplement**

Detailed information on current pricing strategies, competitive landscape, and factors influencing sales outcomes.

- **Launchpad Company Overview**

Foundational document outlining the company's mission, product vision, and market positioning.

- **Module 1 Prompt Pack**

A collection of expertly crafted prompts to guide your AI interactions for effective data synthesis.

DEFAULT PROMPT

Step 4 — Build Your Framework

- Ask Claude to build a pricing & packaging recommendation for Launchpad. Request **3–5 options** with evidence ratings (Strong / Moderate / Emerging), segment fit, and expected effect on churn and expansion. End with a recommended sequence and what to validate first.

If a skill is loaded, run it instead.

Other Pricing Lenses

Packaging Lens (Good-Better-Best)

Design a three-tier packaging structure with target segments, defining features, and the trigger that moves a customer up a tier.

Price-Sensitivity Lens

Map willingness to pay by segment — where you're leaving money on the table and where you're pricing yourself out.

Step 4.5 — Package It as a Skill

Package tonight's workflow into a reusable `willingness-to-pay-signal` skill, then import an expert-built pricing skill and compare what it catches that yours missed. Update your skill accordingly.

Claude With Domain Expertise Sounds Different

When you give Claude a skill built on proven frameworks, it stops giving you research reports — and starts giving you instructions.

Thin prompt

"HERE'S A PRICING RECOMMENDATION."

Useful once. Nothing carries to the next product, the next quarter, or your teammates.

Your skill — willingness-to-pay signal

REPEATABLE, BUT ONLY AS SHARP AS YOU ARE

You package tonight's workflow: the context, the sources to read, the recommendation format. Now it's one prompt on any product's data.

An expert skill — Lenny's pricing-strategy

THINGS A HOMEGROWN PROMPT WON'T SURFACE

Run it on the same data and watch the gap: willingness-to-pay as a measure of value, the ~\$10–15k self-serve ceiling, grandfathering reversible changes, pricing as a living roadmap item.

STEP 5

Pressure-Test & Deliver

Challenge Claude's output before treating it as final, then format it for a real stakeholder meeting.

Pressure-Test Prompts

- What's the weakest link? Where is evidence thinnest?
- Argue AGAINST the #1 opportunity — what would you say?
- What additional data would you want before making a product decision?

Executive Summary Format

Reformat the final framework as a 1-page summary for a VP Product meeting. Include:

- The strategic question
- Top 3 recommendations ranked by evidence strength
- Key risks
- What needs to be validated next

Where Human Judgment Is Required

Weighting Patterns

Claude counts frequency, but frequency $\neq$ importance. One quote from your largest enterprise customer may matter more than 15 Starter-tier tickets.

Reading the Absence

Claude analyzes what's there. You need to notice what's missing — the dog that didn't bark is often the most important signal.

Holding Strategic Tension

The self-serve vs. white-glove tension is real. The "obvious" answer (build self-serve) is complicated by Amanda Greene's perspective. Good synthesis holds both sides.

Validating Emotional Inference

Claude will infer beliefs and feelings from text. If its read doesn't match yours, trust your judgment and push back with specific evidence.

Whose Voice Is Missing?

Churned-and-gone customers can't answer your survey. The quietly dissatisfied don't open tickets. Smaller accounts rarely get interviewed. Claude can only synthesize what's in the room — you have to ask who isn't.

6:20 PM

Hallucination Risks to Watch For

Invented Quotes

Claude may paraphrase or fabricate quotes that sound plausible but weren't in the transcripts. Always ask:
"Can you cite which transcript this comes from?"

False Consensus

If 7 transcripts say one thing and 1 says another, Claude may present the minority as equally weighted. Ask: *"How many data points support this?"*

Hypothesis Confirmation

If you share your hypothesis alongside the data, Claude leans toward confirming it. Challenge it: *"Argue against this using the evidence."*

Generic Recommendations

Claude defaults to broad statements unless you enforce specificity. If a recommendation could apply to any SaaS company, it's too vague — push for Launchpad-specific evidence.

If Your Org Restricts AI Tools

Data sensitivity is a real constraint. Here's how to run this workflow anyway.



Use Synthetic Data in the Lab

The Launchpad dataset is fictional by design. Run the workflow on synthetic data first — get the muscle memory — then apply it to real data within your approved guardrails.



Anonymize or Aggregate

Strip names, company names, and identifiers before uploading transcripts. Aggregate NPS comments by theme rather than pasting raw responses. The synthesis still works — Claude doesn't need to know who said it.



Use an Approved Enterprise Instance

Many orgs have an enterprise Claude or ChatGPT instance with data-handling agreements in place. Check with your IT or legal team — you may already have a compliant path.



Keep Files Local via Cowork

Cowork's desktop app processes files locally before sending context to the model. Transcripts never need to leave your machine. This is often sufficient for orgs with moderate restrictions — verify with your security team.

YOUR MOVE

Make It Yours — Your Data Inventory

The workflow you just ran is only as valuable as the data you run it on. Before you close your laptop tonight, answer one question:

What unsynthesized discovery data will you run this on Monday?

Granola Folder

How many unreviewed call notes are sitting in there? Pick the last 8–10 from a single customer segment and run Step 1 before your next roadmap review.

Zendesk / Support Export

Pull a CSV of last quarter's tickets filtered by your top churn risk segment. That's your data load for Step 2 — ready in 5 minutes.

Last Quarter's NPS

Open-text NPS comments are gold that almost never gets synthesized properly. Drop them into Cowork alongside two interview transcripts and let Step 3 surface the surprises.

The Question That Matters

What decision are you trying to make right now — a roadmap bet, a churn hypothesis, a pricing conversation? Name it. That's your synthesis context for Monday.

✔ You don't need new data. You need to synthesize what you already have.

6:30 PM

Take a Break

BREAK

10 Minutes

Grab a drink, stretch, and come back ready to build.

We'll pick back up at 6:40 — don't wander too far.

6:40 PM

MODULE 2 OF 3

Competitive Intelligence Agents

Stop relying on quarterly decks and Slack hearsay. Build a system that monitors, synthesizes, and informs strategy — continuously.

"Competitive intel is only valuable if it changes decisions."



The Problem Today with Competitive Intel

- ❏ **The shift:** From "what did our competitors do last quarter?" to "here's what changed this week, and here's what it means for our roadmap."

Fragmented signals

Competitive info lives in Slack threads, Intercom chats, G2 reviews, LinkedIn posts, and individual PMs' brains. Nobody has the full picture.

Reactive, not systematic

Most orgs do competitive research when a deal is lost or a competitor launches something big. That's already too late to respond strategically.

Partial data in strategy rooms

Roadmap and pricing decisions get made based on incomplete, stale competitive data. Confidence is high; accuracy is low.

AI changes the game

Continuous monitoring, cross-competitor synthesis, and direct translation into product decisions — without a dedicated competitive analyst.

Cowork today, Python tomorrow

We're building this as a **Cowork Scheduled Task** — not a Python agent running in a terminal. Here's why that matters for your team.

Cowork Scheduled Task

- Runs weekly on its own — set it and walk away
- No terminal, no script to maintain, no waiting on execution
- Review a pre-run brief while it gathers sources in the background
- Anyone on the team can own and iterate it — no Python required
- Connectors (Gong, Zoom, Drive) load internal data automatically

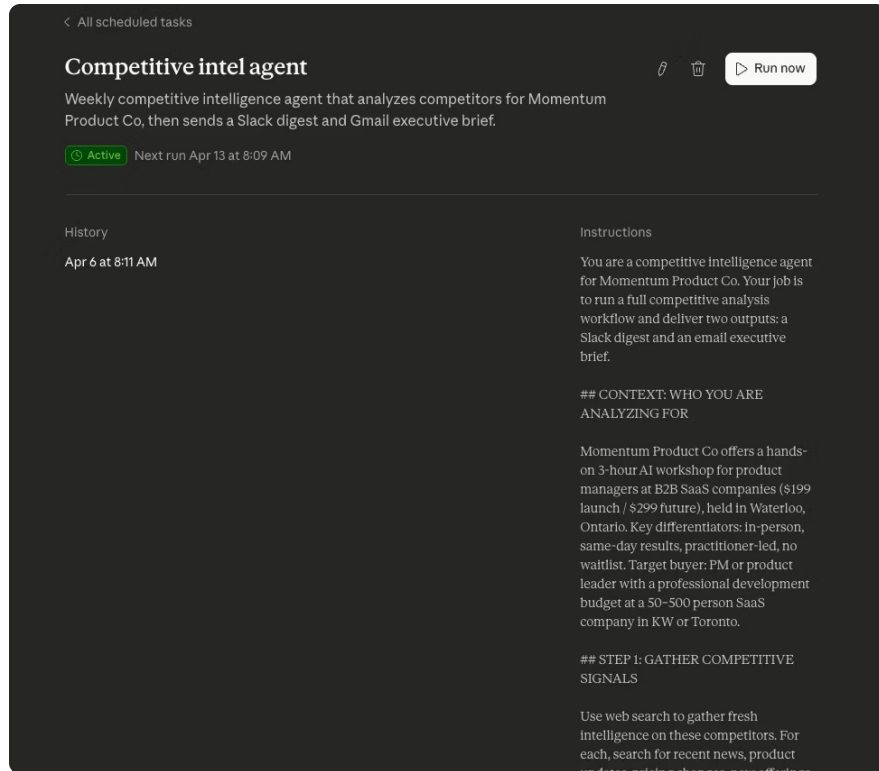
Python Agent (Level-Up)

- Full control over the execution environment
- Scriptable and version-controllable in a repo
- Better for complex multi-step orchestration or custom data sources
- Higher maintenance cost — someone has to own the code
- Build this in Claude Code once you've validated the flow works

- ❏ Start with the Cowork scheduled task. When you've validated the prompts and outputs, the Level-up lane walks you through rebuilding the same flow as a Python agent in Claude Code.

DEMO

Agent Inputs & Outputs



3 min to run, 5 competitors, 3 output formats

Slack Digest

A weekly 3–5 bullet summary of meaningful competitor moves — sent automatically to a #competitive-intel channel. Short, scannable, linkable.

Exec Brief

A structured 1-pager: top competitor moves, strategic implications, recommended responses. Ready for leadership review without extra editing.

Roadmap Inputs


Specific feature gaps, positioning risks, and pricing signals — formatted as direct inputs to your next roadmap or sprint planning session.

The Competitive Intel Workflow

- ❏ The system prompt is the foundation — it tells the agent what you care about, who your customers are, and what decisions the intel should inform. Without it, you get summaries. With it, you get strategic inputs.

⚡ The highest-leverage step

Context in the system prompt. Define your ICP, your differentiators, and the specific decisions this intel should inform. Weak context → generic summaries → zero strategic value.



Synthesis > summaries

Don't just ask "what did Competitor X do?" Ask: *"What pattern is emerging across direct competitors that our roadmap doesn't address yet?"* That's where the value is.

Gather Signals: Internal + External

The scheduled task pulls from two lanes of sources. **Both matter.** External tells you what competitors are saying; internal tells you what they're doing *to your customers*.

External Sources

Competitor Sites & Changelogs

What features shipped? What was renamed or repositioned? Changelog monitoring catches strategic pivots before the press release.

Pricing Pages

Pricing changes are competitive signals. Screenshot and date-stamp them — AI has no reliable sense of when data was collected unless you tell it.

Review Sites & LinkedIn

Capterra, and LinkedIn posts surface what buyers are actually saying — and what competitors are actively promoting to your ICP.

Internal Sources

Webinar & Session Transcripts

Gong and Zoom connectors load these directly into the Cowork scheduled task — no export required. Win/loss calls are especially high-signal.

Sales-Call Notes

Where are reps hearing competitive objections? What are they being told about competitor roadmaps? This is real-time intel your review sites don't have.

Churn Notes That Name a Competitor

"We went with [Competitor X] because..." — these are your highest-signal records. Pull them from your CRM and include them every cycle.

⚠️ WATCH OUT

Where AI Tends to Hallucinate or Overgeneralize

⚠️ Fabricated Specifics

AI may invent specific pricing numbers, feature names, or dates that sound plausible. Always verify any specific claim against the actual source.

📅 False Recency

AI has no reliable sense of when data was collected unless you tell it. It will treat a 2-year-old G2 review as current. Always include collection dates.

📊 Overgeneralization from Thin Data

Three G2 reviews are not a pattern. AI will write confidently about "customer sentiment" based on tiny samples. Label the data volume in your prompt: "Based on 3 reviews, not a representative sample."

🔍 Missing the Absence

AI analyzes what is there. It struggles to identify what is notably absent. Prompt specifically: "What are competitors NOT doing that buyers seem to want?"

How to Validate Outputs Quickly

A 3-question spot check for any AI-generated competitive output

✅ Specificity Check

Does every key claim have a source or quote behind it? Generic observations that could apply to any company are a red flag.

💡 So-What Check

Does the output answer "what should we do differently"? If it's all description and no implication, push Claude for a second pass.

😬 Surprise Check

Does anything in the output genuinely surprise you? If it only confirms what you already knew, either your inputs were thin or the system prompt wasn't specific enough.

If Your Org Restricts AI Tools

Internal competitive data is especially sensitive. Churn notes, sales-call transcripts, and win/loss data can be confidential. Here's how to handle it.

External-Only Mode

Start with public sources only — competitor sites, changelogs, G2 reviews, LinkedIn, pricing pages. You can get real value from the external lane alone while you work out the policy for internal data.

Anonymize Before Loading

Strip company names and rep names from churn notes and sales-call transcripts before loading them. The competitive signal (e.g., "customer left for a competitor offering feature X") survives anonymization.

Use Your Approved Enterprise Instance

If your org has an enterprise Claude or equivalent, run the competitive task there. Internal data stays within your org's data-handling agreement.

Document What You're Doing

Write a one-paragraph summary of the data sources and retention approach you're using. Share with legal or security before operationalizing. Getting sign-off once is easier than asking forgiveness later.

7:20 PM

Take a Break

BREAK

10 Minutes

Grab a drink, stretch, and come back ready to build.

We'll pick back up at 7:30 — don't wander too far.

7:30 PM

MODULE 3 OF 3

Prototyping & Vibe Coding

For product teams who want to go from insight to working prototype — using AI to build, iterate, and validate product ideas **without waiting for engineering or a designer.**



PREREQUISITES

Shipping to Production Today: The Set-up

Module 3 has a short setup. If you haven't done this yet, do it now



Free GitHub Account

Sign up at **github.com** — free tier is all you need. You'll use this to save your build and optionally push to a live URL. QR code on screen.



Free Vercel or v0 Account

Sign up at **vercel.com** or **v0.dev** — free tier covers everything you'll do tonight. This is how you get a live, shareable URL at the end. QR code on screen.



Claude Pro + Cowork Desktop App

Claude Pro (\$28/mo) and the Cowork desktop app installed. We'll be building inside the desktop app — no terminal, no command line required.



AI-powered prototyping changes this.

Product development teams can now go from an insight to a clickable, functional prototype in minutes — something they can put in front of customers or hand to engineering with real specificity.

When to Use This Workflow

Use This When

- You want to test a feature concept before committing eng resources
- You need to *show* stakeholders instead of describing
- You have UX improvements that just need to get done
- You want interactive prototypes instead of static mockups
- You're moving from V1 to V2 and need to visualize what changed

Do Not Use When

- You need production-ready code
- You're building complex backend logic or data pipelines
- You need to work within your existing codebase without eng involvement
- You have less than 20 minutes — even fast prototyping needs planning time

The Vibe Coding Tool Landscape

We're building in Claude tonight. This is for reference — use it when you want to explore further or need UX polish beyond what Claude delivers.

Tool	Known For	Pick This When...
Lovable	Visual editing, click-to-modify, native Supabase integration. Highest UX polish out of the box.	You want to see and tweak what you're building as you go, with a real backend — and UX polish matters to your audience.
Bolt	Full dev environment in the browser — no setup. See and edit the code if you want to. Figma import.	You want the speed of an app builder with the option to look under the hood.
Replit	All-in-one platform: editor, database, hosting, deployment, collaboration. 50+ languages.	You want everything in one place and never want to think about infrastructure.
v0	Made by Vercel. High-quality code output. Git integration with branching and PRs.	You want hosting that scales from prototype to production and a more developer-oriented workflow.

BUILD SMARTER

Plan Before You Prompt

The most common mistake in vibe coding is starting to build before you've thought through what you're building. Claude has a built-in solution: **Plan Mode**.

Enter Plan Mode

Press `shift+tab` before describing your feature. This forces Claude to produce an implementation plan before writing a single line of code.

Review and Push Back

Read the plan critically. Ask: Does this approach make sense? Does it conflict with anything that already exists? Push back before Claude builds — catching bad assumptions here saves 20 minutes of debugging later.

Iterate Deliberately

Describe what's wrong precisely. Vague re-prompts compound issues. "The button doesn't work" is less useful than "clicking Submit doesn't call the API — here's what I expected."

"Product developers who get the most out of these tools are the ones who plan before they prompt. That's your superpower — you already think about user flows, data models, and edge cases."

Give Claude a Design System

By default, Claude prototypes look generic — the same gray cards and blue buttons every time. A markdown design system gives Claude a visual language to adhere to, and designers can own it without touching code.



What a Markdown Design System Is

A plain-text file that defines your colors, typography, spacing rules, component patterns, and tone of voice. Claude reads it at the start of every build and applies it consistently — no design ticket required.



Generate One Tonight

Don't have a design system? Ask Claude to generate one based on your brand colors and a screenshot of your product. You'll have a working markdown file in under 5 minutes. It won't be perfect — but it will make every prototype you build look like *your* product.

- ❏ Load the design system file into Cowork at the start of your session. Claude will reference it automatically throughout the build.

The Designer's Eye

Claude prototypes are functional by default — and generic by default. Knowing when to reach for a different tool is part of the skill.



Use Claude to Validate Cheaply

Claude builds fast and iterates well. For internal prototypes, investor demos, or early customer validation — where you need "does this concept work?" not "is this beautiful?" — Claude is the right tool. Validate the idea before you invest in polish.



Reach for Lovable / Replit When UX Polish Matters

When your prototype is going in front of customers who will judge the experience, or when your brand standard requires visual fidelity, Lovable and Replit produce higher-polish output. The tradeoff is higher cost and longer iteration cycles.



Ask Who's on the Other End

Before picking a tool, ask: who will see this, and what decision will they make?

- Internal engineering alignment → Claude.
- Customer usability test → Lovable.
- Board demo → invest in polish. The answer changes the tool.

7:35 PM

Live Demo: Building in Claude

FACILITATOR-LED · 10 MIN

We're building the **a Workshop Landing Page** — a single build, entirely in Claude. No terminal, no switching tools.

1

Part 1 · 3 min

Enter Plan Mode (shift+tab). Describe the feature and review Claude's implementation plan together.

2

Part 2 · 3 min

Push back on the plan — catch one assumption, change one thing. Then approve it and watch Claude build. Note how it modifies existing code rather than building something new alongside it.

3

Part 3 · 4 min

Iterate: describe one thing that's wrong or missing. Deploy to Vercel from inside Cowork. You'll have a **live URL before your hands-on time starts**.

8:05 PM

DISCUSSION

DEBRIEF

Show, Share & Debrief

Let's see what you built. Share your screen, show your prototype, and we'll debrief together.

Show us what you built

Walk us through your prototype. What does it do? Who is it for?

What surprised you about the process?

Where did the AI take you somewhere unexpected — good or bad?

Where did you have to fight the tool?

What broke, confused you, or required the most iteration to get right?

Would you show this to a customer or engineer tomorrow?

What's missing? What would you add in the next 30 minutes if you had them?

The Vibe Doom Loop

"The tool makes assumptions about how to build your thing — the folder structure, the component hierarchy, the data model. You can't see those assumptions. When you add something later that contradicts them, things break. And then you prompt to fix the break, which creates new assumptions, which break something else."

The fix isn't better prompting. It's planning before you prompt. Think through: What data does this feature need? How does it connect to what already exists? What are the edge cases? That's product management — you already think this way about products you ship.

Plan Mode

`shift+tab` forces an implementation plan before any code is written

Review & Push Back

Critique the plan before Claude builds — catch bad assumptions early

Iterate Deliberately

Describe what's wrong precisely; avoid vague re-prompts that compound issues

The AI Code Review Loop

Once the plan is ready, start a fresh conversation — this resets the context window and helps avoid context rot. Ask the new agent to implement the plan.

If planning went well, the agent should be able to work on its own. When done, its inclination is to ask you to test the changes — but don't do this yet. Instead, have an agent check the coding agent's work.

What the AI Code Reviewer Does

- Reads the plan, scans the codebase, and evaluates the changes
- Looks for bugs, duplicate code, and new patterns or technologies that are not needed
- Looks for over-engineering
- Primary goal: help avoid the vibe coding doom loop
- Makes sure the coding agent follows good software engineering practices

Three Areas of Extra Focus

Error Handling

A critical area often overlooked in vibe coding

Test Coverage

Ensures changes are verifiable and regression-safe

Security

Catches vulnerabilities introduced by AI-generated code

Prototype → PR, Not PRD → AI-Native Shipping

Tonight you've been working on one end of a spectrum. Here's where it's heading — and where most orgs actually are.

1

PRD → Design → Dev → Ship

Traditional product process. High coordination cost. Works — but slow. Most orgs still live here for most features.

2

Prototype → Validate → Handoff

What you built tonight. Faster signal. Lower cost to validate. Still hands off to engineering for production. **This is the highest-leverage move for most PMs right now.**

3

Prototype → PR → Ship

PM-built prototype becomes a PR. Engineering reviews, merges, and ships. Requires eng buy-in and a codebase Claude can read. Some teams are here today.

4

AI-Native Shipping

Agents write, review, and ship code with minimal human intervention. Most orgs aren't here yet — and that's fine. The journey is the value.

- ☐ Most orgs aren't AI-native yet. That's the gap Momentum helps close — as senior advisors, fractional leaders, and through the **Build with AI 4-Week Maven Cohort**. Ask us about it after the session.

8:15 PM

WHAT'S NEXT

Most teams don't have an AI workflow problem — they have an operating problem

What you've done

You've tried the tools

You've built a few workflows

You've seen some speed

But...

Insights aren't trusted

Outputs feel shallow

Work isn't repeatable

It's not changing product decisions



AI is being used — but not integrated

What it looks like to actually operationalize AI in product



Turn workflows into repeatable systems



Move from outputs → decision-quality insights



Build shared systems across your team and organization



Go from prototype → production-ready thinking



Design your AI product operating model

"If you're already experimenting, this is how you make it real."



Build with AI: 3-Week Cohort

STARTS JUNE 15 · ONLINE · \$599

\$599

25% off COURSE ATTENDEES

- 3 Wednesdays mornings live (9-10:30 am)— 1.5 hours education
- Optional project each week
- 1 hour Friday live Q&A session (12-1 pm)
- Small cohort, real work, real outputs
- 3 lessons, 3 projects, 8 hours live education
- Momentum Product Community Access



Signal to Strategy: How PM's Accelerate Decision-Making with AI

Heather Gawel

Ex-VP Product | AI Advisor to PMs

Babajide Okusanya

AI Engineer | 2,000+ trained

[View Syllabus](#)



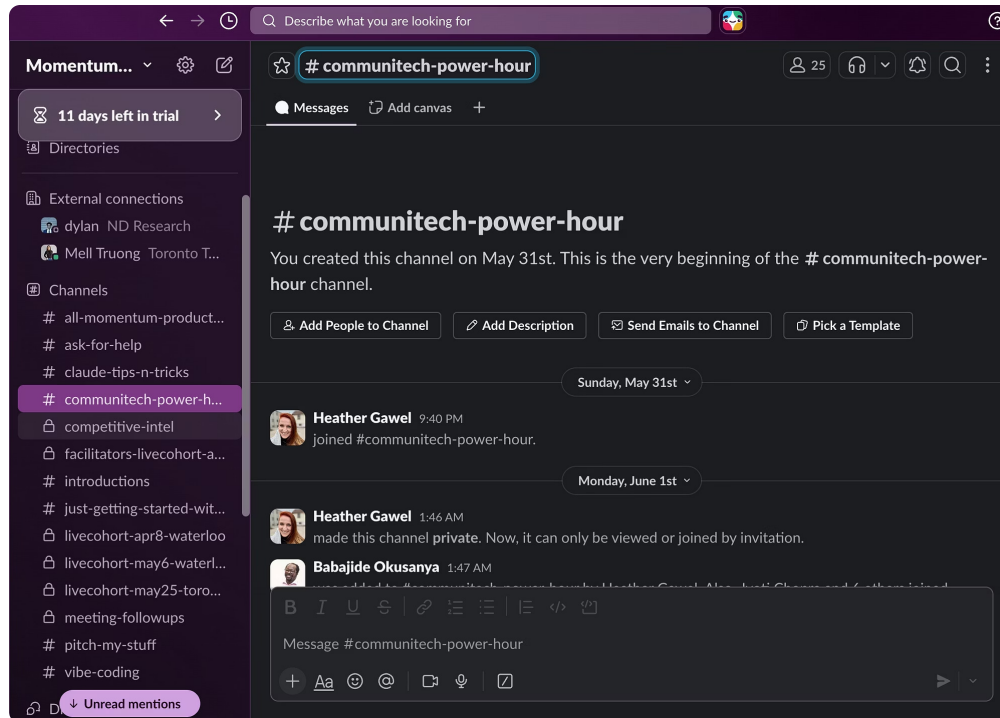
Most "AI for product" courses teach you how to build with AI. This one teaches you how to synthesize customer inputs, build agents and prototypes that surface signal, and ultimately decide product direction with conviction.

AI doesn't just accelerate delivery — it compresses the cycle between signal and strategy. The product leaders pulling ahead right now aren't the ones who ship fastest. They're the ones whose roadmaps are grounded in customer truth, competitive reality, and live product signal — and who can prove it.

Over four weeks, you'll build a decision system powered by AI: one that ingests customer insight, stress-tests it against your strategy, validates ideas before you build, integrates with existing engineering workflows and closes the loop after launch.

BEFORE WE START

Join Momentum AI Working Group on Slack



Same QR from your pre-work email

Download the working Files

How'd We Do?

We'd love your feedback — it takes 2 minutes and helps us make the next session even better.

If you enjoyed this course and would recommend us, please share a testimonial. This has the greatest impact on our ability to reach more people.

[Share Your Feedback →](#)

Thanks for building with us tonight. 🚀



Thank You 🙌

Tonight you didn't just learn about AI — you built with it. Big Thank You to Communitel for partnering with us.

Take what you built tonight and run it on your own data this week. That's where the real learning happens.

 **Customer Synthesis** Turn research into decisions

 **Competitive Intel** Monitor your market automatically

 **Working Prototype** From idea to clickable

Stay connected:

Momentum Product Company

momentumproductco.com