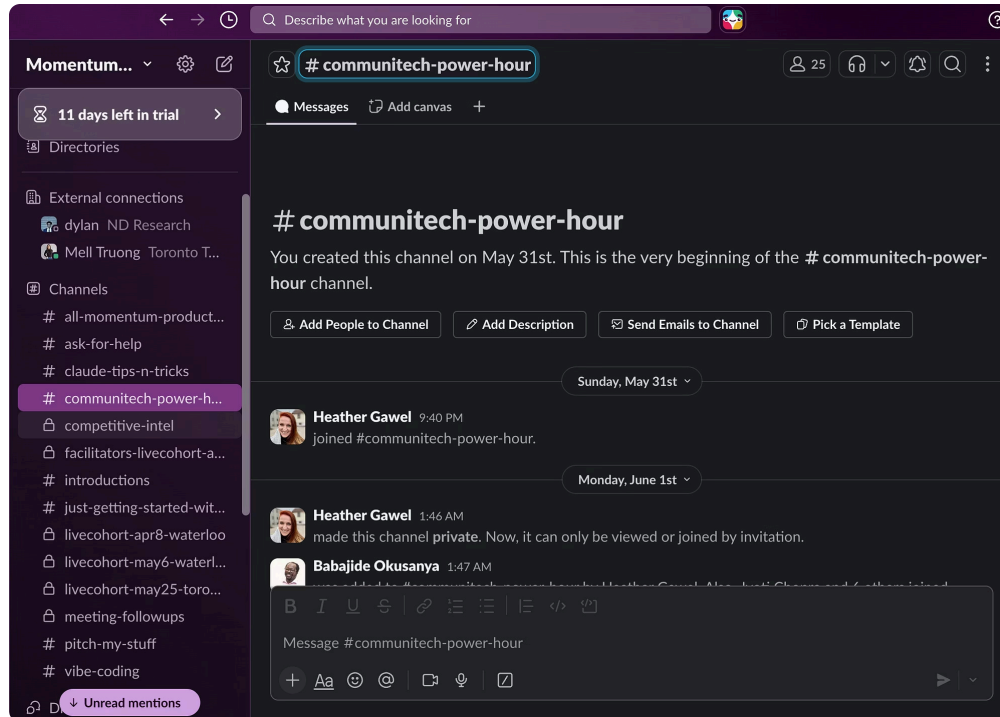


BEFORE WE START

Join Momentum AI Working Group on Slack



Same QR from your pre-work email

Download the working Files

COMMUNITECH × MOMENTUM PRODUCT CO.

AI POWER HOUR

Turn Customer Feedback Into a Decision

Session Details

June 8, 2026 · 5:30–6:30 PM

Communitech, Waterloo Region

Hosted by Communitech × Momentum
Product Co.

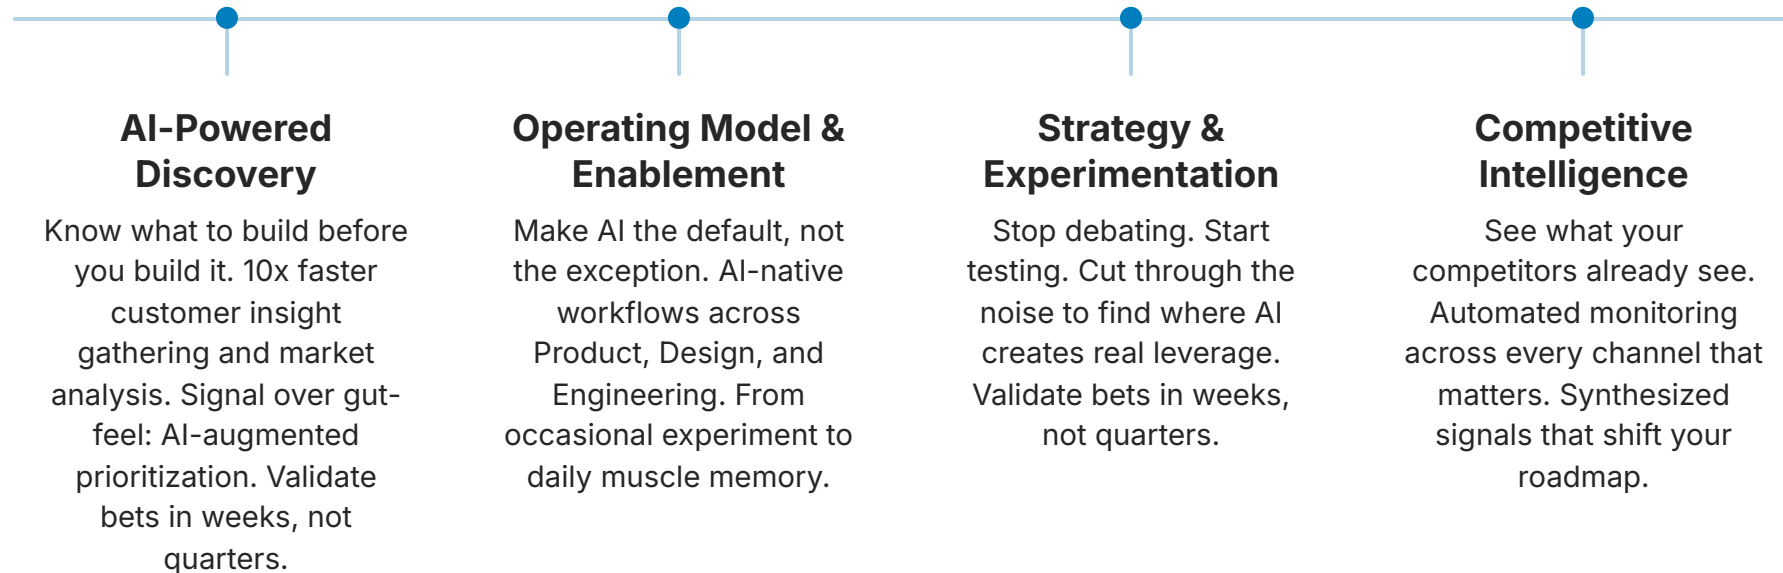


Momentum Product Co

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.

Where We Add Value



📄 We specialize in **bounded engagements** to upskill your existing team — so they don't need us.

👤 YOUR INSTRUCTOR

Heather Gawel

Founder, Momentum Product Company, Product Practice Lead [Waterloo/Toronto]

📄 About Me

I have lived these workflows and been on the forefront of implementing AI-first processes across product, design, and engineering — and have the battle scars to prove it.

I combine **strategic product leadership** with hands-on AI enablement to transform how teams discover, prioritize, and validate.

I'm committed to upskilling teams and organizations to capitalize on this technology — not with theory, but with **real workflows that stick**.

I am excited to build with you and welcome you to the Momentum Product Community.



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

EXERCISE

Discovery & Customer Insight Synthesis

A workshop guide for cross-functional teams who want to turn raw customer conversations, support tickets, and survey data into patterns, opportunities, and product decisions — using AI as a research partner, not a note-taker.

This adds value when overlaid onto strategy, roadmaps and decision frameworks. More data, without context, is useless.



The Problem This Workflow Solves

2–3 Days Lost

The average PM spends 2–3 days synthesizing a single round of customer discovery — time that rarely produces proportional value.

Data Silos

Interview notes sit unreviewed in Granola. Support tickets pile up in Zendesk. NPS comments get skimmed once a quarter.

Stale Insights

By the time insights reach a roadmap discussion, they're either outdated or stripped to a single bullet that loses all nuance.

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**.



When to Use This Workflow

✓ Use This When

- You have 5+ customer calls and need to find patterns quickly
- You're preparing for a roadmap review or strategy discussion
- You want to combine interviews, tickets, and surveys into one coherent picture
- You have a backlog of discovery calls nobody has synthesized

⊘ Do Not Use When

- Fewer than 5 calls — patterns won't be meaningful
- You need statistically significant results — this is qualitative synthesis
- You're under 20 minutes — start with fewer sources
- Interviews covered very different user types — segment first

What You'll Need & The Sample Data

About the Sample Data

The workshop uses data from **Launchpad** — a fictional B2B SaaS onboarding platform. Your role: reduce annual logo churn from **14% to under 10%**.

Required Inputs

- Interview transcripts or notes from at least 5 customer calls (Granola, Gong, Zoom, or any format)
- Support tickets or customer feedback (CSV, spreadsheet, or raw text)
- NPS or survey data with open-text comments
- A clear decision lens — what specific question are you answering?
- Optional: existing hypotheses you want to challenge

The dataset includes:

- 10 discovery interview transcripts
- 75 support tickets from Q1 2026
- 120 NPS responses with open-text comments
- A company overview with competitive landscape and churn analysis

Why Cowork OR Claude.ai for This Module

Exercise runs entirely inside **Cowork** — Claude's desktop app OR Claude.ai. 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.

From Synthesis to Signal

This module isn't just for PMs. Whoever is in the room, the goal is the same: **validate assumptions instead of confirming them.**

For Product & Engineering

This is where product, and eng decide *what to build*. Buried signal — the friction patterns, the unspoken workarounds — surfaces here, not in a quarterly deck. Use the output to gut-check your next roadmap cycle before you commit to a sprint.

For Partnerships

"Here's everyone who asked for this integration — and the churn we lost when we didn't build it." Synthesis gives you the receipts for partnership conversations that used to rely on anecdote.

For Design

Raw transcripts carry emotional texture that ticket counts don't. Use the output to anchor journey maps and How-Might-We statements in what customers actually said — not what the team remembers.

Human Judgment Always

Claude surfaces signal. You decide what it means. Frequency isn't importance. Context isn't in the data. The synthesis is the starting point for a decision — not the decision itself.

Two Lanes

Pick the path that matches where you are. Both lanes produce the same core artifact — the synthesis framework. The Level-up lane goes further.

Core Lane

Guided, prompts provided

- Follow the participant guide step-by-step
- Prompts are written for you — focus on reading the output and pushing back
- Upload sample data via Cowork desktop app
- Deliverable: a completed Opportunity Solution Tree on Launchpad data

Level-Up Lane

Pull live data, modify the skill, go further

- Connect your own Gong, Zoom, or Drive data via Cowork connectors
- Modify or write the synthesis skill to match your ICP and decision lens
- Have Claude script the frequency counts — how many transcripts, tickets, and NPS comments mention each theme
- Deliverable: a synthesis on your own real customer data, ready to share Monday

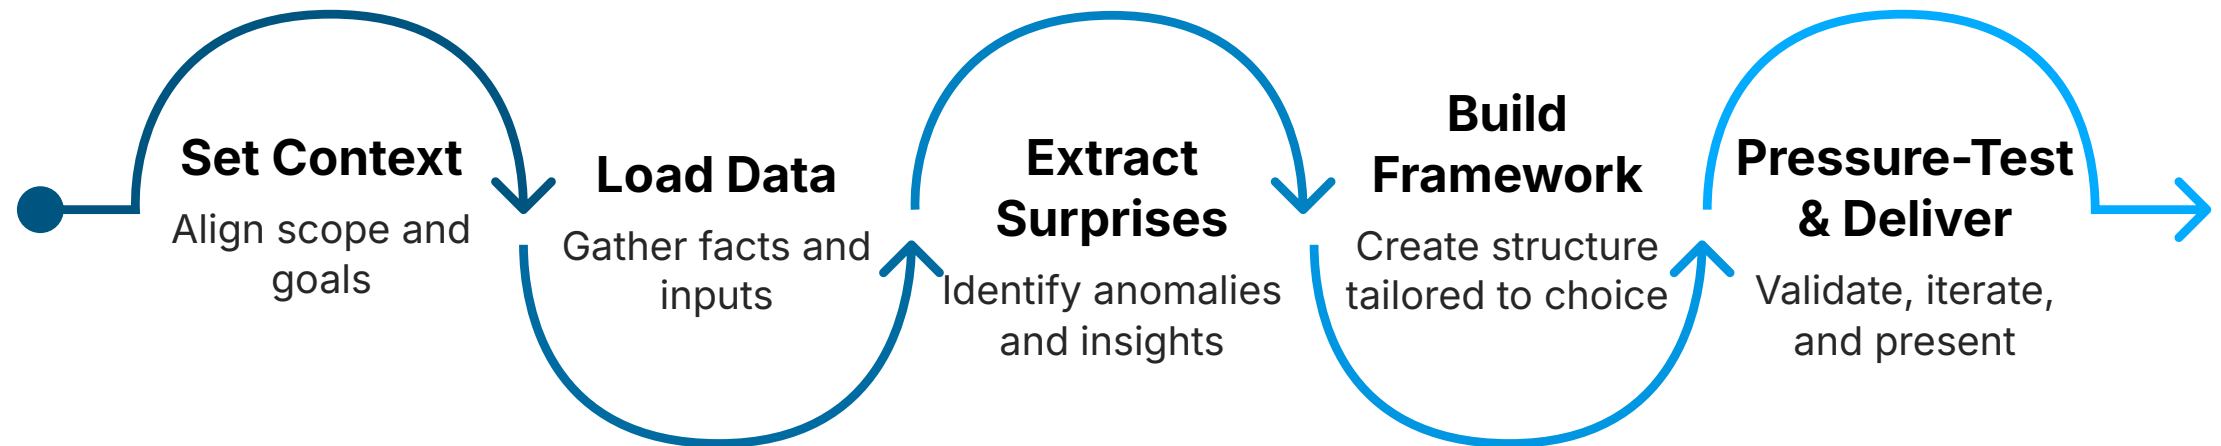
Choose Your Synthesis Framework

The **Opportunity Solution Tree** is the default for this workshop — it produces the most decision-ready output in 40 minutes and works for every role in the room. Jobs-to-Be-Done is the recommended alternate. Designers: see the Designer Track below.

Framework	Best For	Deliverable	Difficulty
1. Opportunity Solution Tree ★ DEFAULT	Mapping pain to product opportunities — for the whole cross-functional room	Visual tree: outcomes → opportunities → solutions	Medium
2. Jobs-to-Be-Done ✂ Alternate	Understanding WHY customers hire/fire your product	Job statements with outcome expectations	Medium
Other Lenses / Appendix	See participant guide for: Problem Prioritization Matrix, Hypothesis-Driven Product Brief, Customer Segmentation + Needs Map	Varies	Easy–Hard

Use the Participant Guide: [01.Participant_Guide_Discovery_Launchpad_FINAL](#) FOLLOW ME or MOVE AHEAD

The 5-Step Workflow



Steps 1–3 are identical regardless of framework. Steps 4–5 diverge based on your chosen framework. The entire workflow is designed to complete in under 40 minutes.

Steps 1–3: Context, Data & Surprises

1

Set Your Synthesis Context

Before sharing any data, establish business context and your specific question. Tell Claude the churn target, data sources, and what you want it to do — identify cross-source patterns, contradictions, absences, and high-impact vs. high-frequency issues. **Skipping this step produces generic summaries, not strategic insights.**

2

Load the Data

Start with 4–5 interview transcripts (not all 10 — leave room for follow-up), then add the Support Tickets CSV and NPS Responses CSV. Always orient Claude on column meanings. Good output looks like: *"The mobile experience issue appears in 3 interviews, 10 tickets, and 12 NPS comments."*

3

Extract Surprises Before the Framework

Ask Claude what surprised it — patterns only visible across sources, where segments disagree, what's notably absent, and the single most important insight. Push back if it returns high-frequency themes disguised as surprises: *"That's a pattern, not a surprise."*

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.

Without a Skill

📄 "Here's what I found, here's what it means, here are your options."

Structured. Helpful. But shaped like a summary — not a mandate.

With a Skill

"The skill is really insistent that the point of feedback analysis is action, not understanding. Geoff Charles's 'every ticket is a product failure' is not an epistemological claim — it's an instruction: share every ticket with the relevant PM monthly. Shaun Clowes's 'feedback river' isn't a metaphor, it's a process mandate."

Domain-grounded. Opinionated. Tells you what to do, not just what to think.

📄 "The difference isn't intelligence — it's context. Skills give Claude the frameworks to move from insight to action."

Step 4: Build Your Framework



Opportunity Solution Tree ★ Default

Map desired outcome → 4–6 opportunities → 2–3 solutions each. Rate evidence strength (Strong/Moderate/Emerging), note segment relevance, and include compelling quotes. Organize from strongest to weakest evidence.



Jobs-to-Be-Done (Alternate)

Write job statements: *"When [situation], I want to [motivation], so I can [outcome]."* Rate how well Launchpad delivers each outcome (Strong/Adequate/Failing) and identify which underserved outcomes drive churn.



Designer Track

Generate How-Might-We statements anchored in the top opportunities. Then ask Claude to build a current-state journey map — key stages, emotional highs/lows, and the moments most correlated with churn. Use as input to your next design sprint.

- ❏ Other lenses (Prioritization Matrix, Hypothesis Brief, Segmentation Map) are available in the Appendix of your participant guide for use after the session.

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

Synthesis vs. Summary

The most important distinction in this exercise is between **summary** (restating what each source said) and **synthesis** (finding what the sources collectively reveal — including contradictions and absences). Use this corrective prompt whenever Claude produces summaries disguised as synthesis.

"That's a summary of each source, not a synthesis across them. Stop listing what each interview said. Instead: what do these transcripts, tickets, and NPS comments collectively reveal that no single source shows on its own? What pattern only becomes visible when you cross-reference?"

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.

Discussion Prompts for Debrief

What did Claude surface in Step 3 that surprised you?

What does that tell you about what you would have missed if you'd synthesized manually?

Did anyone get different results using different frameworks?

How did the framework shape what you noticed — and what it caused you to overlook?

Where did you have to push back on Claude?

What prompt change produced better output? What would you do differently next time?

If you ran this on your own customer data tomorrow...

What's the first question you'd ask? What data do you already have that you haven't synthesized?

Key Takeaways



Context First, Always

The synthesis context you set in Step 1 determines everything. Generic input produces generic output. Anchor Claude to your specific business goal before sharing a single data point.



You Are the Final Judge

Claude surfaces patterns; you weigh them. Frequency isn't importance. Strategic tensions require human judgment. Always pressure-test before bringing output to a stakeholder.



Demand Cross-Source Synthesis

Push Claude beyond single-source summaries. The most valuable insights only emerge when interviews, tickets, and NPS data are cross-referenced — including contradictions and absences.



Under 40 Minutes

What used to take 2–3 days can now be done in a focused 40-minute session — leaving more time for the decisions that actually require your expertise. This workflow can be run weekly on new interviews, new NPS data, etc.

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.

Upcoming Events FREE Communitech Events

Continue your AI journey with these specialized workshops and labs.

AI Power Hours



Prompts That Push Back

 **June 15, 2026 | 5:30–6:30 p.m.**

Master the “push back” technique that helps AI critique and improve its own responses instead of confidently guessing.

 [Register](#)




What’s an Agent, and How Do I Build One?

 **June 22, 2026 | 5:30–6:30 p.m.**




Build a simple browser-based AI agent that can complete real tasks for you.

 [Register](#)



Advanced AI Application Workshop

Ready to move beyond prototyping? This workshop goes deeper on the workflows already changing the job — turning research, docs, and support tickets into decisions in minutes, and building AI agents that actually run the work instead of just demoing it.

-  **June 18th, 2026**
-  **5:30 PM - 8:30 PM EDT**
-  **Communitech**

Scan the QR code to register and secure your spot!



Or register directly at luma.com/O8m5ihzd

MOMENTUM PRESENTS



Build with AI: Tech Worker Edition

A hands-on lab for the tech workers building with AI today.

Thursday, June 18, 2026

5:30 – 8:30 PM

SPONSORED BY

Communitech



WHAT'S NEXT

ONLINE → 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 m) — 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.

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. 🚀

