



The Vibe Lab

Session 2 • Attendee Workbook

Everyday AI Use — an interactive vibe-coding session

Thursday, September 24, 2026

KRESA Career Connect Campus

THIS WORKBOOK BELONGS TO

You are going to write prompts today

Not watch someone else do it. This workbook is yours to write in, mark up, disagree with, and take home.

Everyone in this room is at a different place with AI. Some of you were here in June. Some of you opened an AI tool for the first time this week. Both are fine. The lab is built so that you work at your own pace and compare notes with the person next to you.

The goal is not to make you an expert in two hours. The goal is that you leave with **one repeatable process** you can run on your own work on Friday morning.

How to use this workbook

- **During the session** — the lab pages follow each step in order. Every prompt is printed so you never have to squint at the screen.
- **The blank lines are the point.** Write down what your AI actually gave back. Comparing your results to your neighbour's is where the learning happens.
- **After the session** — the take-home pages are yours to keep: practice prompts, safety guidance, a tool comparison, a glossary, and a 30-day plan.

A note on pace

You will not finish every page today, and you are not meant to. Work through the lab with the room, then take the rest home. Nothing in here expires.

WHAT I WANT TO WALK OUT OF HERE ABLE TO DO

What is in here

Your day at a glance	4
Getting set up	5
The big idea	6
Five habits that make results repeatable	7
What “be specific” actually buys you	8
The four-step loop	9
Anatomy of a good prompt	10
Five follow-ups worth memorizing	11
Today’s mission	12
The lab: Steps 1–4	13
Debrief	19
Using AI safely at work	20
Trust, but verify	21
Set your own ground rules	22
Practice prompts for real work	23
Build your own prompt (worksheets)	26
Choosing a tool	28
Glossary	29
Your next 30 days	31
Notes	32

Your day at a glance

1

Introductions and polling questions

Meet the room. Name, company and role, and where you are with AI.

2

Lab overview

Tom Minto walks through the dataset and the two paths you can pick from.

3

Getting repeatable results

The habits and the four-step loop that make AI useful instead of random.

4

Interactive lab

You write the prompts. We compare what everyone got back.

What to bring

- ☐ A laptop — this is hands-on, you will be typing the whole time
- ☐ A login for any major AI tool (free is fine — see the next page)
- ☐ The two files we emailed you, downloaded and easy to find
- ☐ One real task from your own work that you wish took less time

THAT ONE REAL TASK — WRITE IT HERE

ONE THING I WANT TO ASK TODAY

Getting set up

No AI account yet? You can be ready in about two minutes, and it does not cost anything.

Get Claude — free (used in today's demo)

- Go to **claude.ai**
- Sign up with an email address or a Google account
- No credit card, no trial expiration
- The free plan is plenty for everything we do today

Already use something else?

- **ChatGPT** (chatgpt.com) and **Grok** (grok.com) both have free tiers
- Every prompt in this workbook works in any of them
- Only the final “save it as a skill” idea is specific to Claude

Heads up: free plans limit how many messages you can send in a few-hour window. That is more than enough for this lab, but if you hit it, pair up with a neighbor and keep going.

WHICH TOOL ARE YOU USING TODAY, AND WHY?

The big idea

Treat AI like a new teammate

It is capable, it is fast, and it has **zero context** about you, your company, or what a good result looks like in your world. Every bad AI experience you have had traces back to that gap.

Think of it as a brilliant intern on their first morning. You would not walk past their desk, say “fix the spreadsheet,” and walk away. You would tell them what the spreadsheet is, what “fixed” means, and what to do if something looks wrong.

Your job is the onboarding. That is the whole skill.

***Repeatable means:** you can do the same thing next week, on different data, and get the same quality — without rebuilding the whole conversation from scratch.*

What usually goes wrong

- One vague line, hit enter, disappointing answer, conclude “AI is overhyped”
- Letting it change everything at once, then not being able to tell what it did
- Getting a great result and never writing down how you got it

WHERE HAVE YOU SEEN THIS HAPPEN TO YOU?

Five habits that make results repeatable

- 1 Give it a role and a goal.** Tell it who it is helping and what a good result looks like. “You are helping a bookkeeper close the month” beats no framing at all.
- 2 Be specific.** Vague in, vague out. Name exactly what you want it to check. This is the habit that pays off most today.
- 3 Ground it in your data.** Attach the file. Do not make it guess from memory, and do not paste a summary when you can give it the real thing.
- 4 Review before you act.** See what it found before you let it change anything. Beginners let the AI “fix” everything at once and then cannot tell what changed.
- 5 Capture what worked.** Save the process so next time is faster. This is the step almost everyone skips.

WHICH HABIT DO YOU ALREADY DO? WHICH ONE WILL YOU ADD?

What “be specific” actually buys you

Same file. Same AI. Two prompts. Wildly different results.

Vague

Check this data and let me know if anything looks wrong.

Typical result: it eyeballs the categories, finds two or three obvious things, and tells you it looks “generally consistent.”

Specific

For each row, check that the category fits the description, and that the total equals quantity × rate.

Also flag: duplicates, blank fields, negative or impossible values, end dates before start dates, and any date outside 2025–2026.

Typical result: a numbered list with the exact record IDs, what is wrong, and what it should be.

The lesson: the AI is not holding back. It answers the question you asked. A vague question is a small question.

The four-step loop

One pattern that works for almost any task involving a file — spreadsheets, invoices, meeting notes, contracts, survey results.

- 1 **Review** — Look, do not touch. “Tell me what you see. Do not change anything yet.”
- 2 **Validate** — Check the rules. “Here is what correct looks like. Show me everything that breaks it.”
- 3 **Act** — Now do one thing. Clean it, summarize it, chart it, or draft from it.
- 4 **Capture** — Write down the process so you can rerun it next month in a fresh chat.

***Training tip:** start today by telling the AI [not](#) to save anything to memory, so practice data does not follow you into real work. It is also useful to know that [toggle](#) exists.*

Where else could this loop run?

List three tasks at your job that involve opening a file and checking it.

Anatomy of a good prompt

Not a magic formula. Just the five things that are usually missing.

R

Role & goal — who it is helping, and what a good result looks like.

“You are helping an office manager get ready for a board meeting.”

C

Context — the file, the background, the constraints it cannot guess.

“We are a 14-person HVAC company in Kalamazoo. Our fiscal year ends in June.”

T

Task — one clear instruction, with the specifics named out loud.

“List every invoice where the total does not match quantity × rate.”

F

Format — how you want it back.

“A table with invoice number, what is wrong, and what it should be.”

G

Guardrails — what not to do.

“Do not change anything yet. Do not guess — if you are unsure, say so.”

Rule of thumb: if you would have to explain it to a new hire, you have to say it to the AI. If a new hire would not need it, skip it.

Five follow-ups worth memorizing

The first answer is a draft, not a verdict. These five replies improve almost any result.

“Show your work.”

Ask how it reached each conclusion. Wrong reasoning with a right answer is still a problem.

“What did you have to assume?”

Surfaces the guesses it quietly made about your business.

“Rate your confidence on each item.”

Gives you a triage order for what to check by hand.

“What did I not ask you to check that you would recommend?”

Often the most valuable sentence in the whole conversation.

“Do that again, but shorter / for a customer / as a table.”

Reformatting is nearly free. Ask for the version you actually need.

WHICH FOLLOW-UP GOT YOU THE MOST USEFUL ANSWER TODAY?

Today's mission

Find what is wrong in this service log — without hunting for it by hand.

KalFlow_Service_Tickets_2025-2026.csv

12 months of IT service tickets • 161 tickets • July 2025 – June 2026

Columns: TicketID • DateOpened • DateClosed • Customer • Technician • ServiceType • Description • Priority • Hours • Rate • BilledAmount • Status

Some tickets are miscoded, some are mis-billed, and some have dates that are simply impossible. All of it was planted on purpose. Your job is to get the AI to find it, then either clean it up or report on it.

New this session: *several errors only show up when you ask the AI to compare **two columns against each other** — is the billed amount really hours × rate? Did the ticket close before it opened? Nobody finds those by scrolling.*

Pick your path

A

Operations Support — produce a cleaned, flagged copy a human can approve.

B

Reporting & Insight — build charts and show how the errors would have distorted them.

MY PATH TODAY, AND WHY

BEFORE I START — WHAT DO I EXPECT THE AI WILL MISS?

Review

Get the AI oriented before it changes anything.

Prompt — adapt the wording to sound like you

Please don't save anything to memory – this is a training exercise.

I've attached a log of IT service tickets. Review it and tell me: the columns, the date range, the service types and statuses used, and anything that looks inconsistent or off.

Don't change anything yet – just summarize what you see.

Three things are baked into that prompt on purpose: the **no-memory instruction**, the **role and goal**, and the **guardrail** at the end. Notice how much work that last line does.

WHAT DID YOUR AI NOTICE ON THE FIRST PASS?

Bonus: did it spot the two different spellings of one service type before you even asked it to look for errors?

Validate

This is where “be specific” earns its keep.

Prompt

Now validate this data. For each ticket, check that the `ServiceType` fits the `Description`, and that `BilledAmount` equals `Hours × Rate`.

List every ticket that looks wrong – show the `TicketID`, what's wrong, what it should be, and why.

Also flag: duplicates, blank fields, negative or impossible values, close dates before open dates, and invalid or out-of-range dates.

Name these out loud — check them off as your AI catches them

- ☐ Does the `ServiceType` match what the `Description` actually says?
- ☐ Does `BilledAmount` equal `Hours × Rate`?
- ☐ Two spellings of the same service type, status, or customer?
- ☐ Anything blank that should not be?
- ☐ Negative amounts, zero rates, or impossible hours?
- ☐ A close date before its open date?
- ☐ Any date outside July 2025 – June 2026, or invalid entirely?
- ☐ Exact duplicate rows under different ticket IDs?

What did yours catch?

Write down the ticket IDs your AI flagged. Then compare with the person next to you — slightly different prompts produce noticeably different lists.

TICKETS MY AI FLAGGED

WHAT MY NEIGHBOUR CAUGHT THAT I MISSED

Act — pick one

Path A — Operations Support

Prompt

Create a cleaned copy with a "Suggested Fix" column and a "Flag" column noting each issue.

Don't delete anything – just mark and suggest, so a human can approve.

Notice we ask it to suggest, not silently overwrite. Human stays in the loop.

Path B — Reporting & Insight

Prompt

Build charts: billed revenue by service type, and monthly ticket volume over the year.

Then tell me how the errors you found would distort these charts if we hadn't caught them.

Watch what one 98-hour ticket and one duplicate do to an otherwise sensible chart.

What came back?

WHAT THE AI PRODUCED, AND WHETHER YOU WOULD TRUST IT

WHAT YOU WOULD CHANGE ABOUT YOUR PROMPT IF YOU RAN IT AGAIN

Capture

The step people skip, and the one that pays you back.

Prompt

Write a short Markdown file documenting the process we just followed, so I can reuse it in a new chat.

Include: the goal, the step-by-step prompts, what to check for, and the valid service types and statuses.

Keep it concise.

Save that file somewhere you will find it again. Next month: new chat, paste the file, attach new data, and you are instantly back to today's quality with none of the re-explaining.

Run it three or four times, refine it as you go, and it becomes a **saved skill** — a one-click version of this whole routine. We are not building that today. Today you take step one.

Do this before Friday: open a new chat, paste your process file, and run it on something real from your own work. That second run is what makes it stick.

WHERE I SAVED MY PROCESS FILE

Debrief

What did your AI catch — and miss?

Which errors needed you to compare two columns?

Where could you run this loop at work this week?

The takeaway: the difference between a great result and a mediocre one was almost always how specific the prompt was — not which tool you used.

Using AI safely at work

Plain-language guidance. Not legal advice, and not a substitute for whatever policy your employer has — check that first.

Think of a prompt like an email to an outside vendor

You would not paste a customer's social security number into an email to a vendor you just met. Same instinct applies here. The question is not "is AI safe," it is "would I be comfortable if this text left the building?"

Keep out of the chat box

- ☐ Social security numbers, dates of birth, driver's license numbers
- ☐ Full credit card, bank account, or routing numbers
- ☐ Patient health information, or anything covered by HIPAA
- ☐ Passwords, API keys, and login credentials — ever, for any reason
- ☐ Employee records: reviews, salaries, discipline, medical leave
- ☐ Signed contracts and terms you are not free to share
- ☐ Anything a customer gave you in confidence

Generally fine

- ☐ Your own drafts, notes, and half-finished thinking
- ☐ Public information, and anything already on your website
- ☐ Data you have stripped of names and identifying numbers
- ☐ Made-up or sample data that behaves like the real thing

When in doubt: anonymise first. Replace real names with Customer A, Customer B. The AI does not need to know who they are to check your math.

Trust, but verify

It can be confidently wrong

AI tools generate plausible text. Most of the time plausible and correct are the same thing. Sometimes they are not, and the wrong answer arrives in exactly the same confident tone as the right one. There is no tell.

Always check by hand

- ☐ Numbers that will end up in an invoice, a filing, or a board packet
- ☐ Names, dates, dollar figures, and anything quoted as fact
- ☐ Citations, statistics, and links — open them
- ☐ Legal, tax, medical, or HR conclusions — get a human professional
- ☐ Anything going to a customer under your name

Memory and history

Most tools remember your conversations, and some build a running profile of you. That is useful at work and unhelpful during a training exercise. Learn where the memory setting lives in the tool you use, and know how to say “*do not save this.*”

Keep a human in the loop

The pattern that holds up: **AI drafts, a person decides.** Ask it to flag and suggest rather than silently change. Anything that touches money, people, or a customer gets human eyes before it goes out.

One more: AI is not confidential in the legal sense. If a conversation would be awkward in a deposition, do not have it in a chat window.

Set your own ground rules

Five minutes of thinking now saves an uncomfortable conversation later. Fill this in for yourself or your team.

What we will use AI for

What we will not use AI for

What never goes into a chat window

Who reviews AI-assisted work before it leaves the building

Do we tell customers when AI helped? If so, how?

Practice prompts: writing & email

Twelve prompts across three pages. Adapt the wording, attach your own file, and change the details in brackets.

1. Turn a mess of notes into a clear email

"Here are my rough notes from a call with [customer]. Write a short follow-up email confirming what we agreed, what I owe them, and by when. Keep it under 150 words and sound like a person, not a press release."

2. Say no without burning the relationship

"I need to tell [customer] we cannot hit the date they asked for. Draft three versions: one warm, one straightforward, one brief. Do not over-apologise, and give them a real alternative."

3. Tighten something you already wrote

"Cut this by a third without losing anything that matters. Tell me what you cut and why, so I can push back."

4. Pressure-test before you send

"Read this as if you were the customer receiving it. What would confuse you, annoy you, or make you reply with a question I should have answered already?"

WHICH ONE WORKED? WHAT DID YOU CHANGE?

Practice prompts: data & spreadsheets

5. The health check

“Review this spreadsheet and tell me: the columns, the date range, and anything inconsistent. Do not change anything yet.” — *the Review step, on your own data.*

6. The cross-field audit

“Check that [total] equals [quantity] × [rate] on every row, and that every [end date] falls after its [start date]. List every row that fails, with the ID and the correct value.”

7. Find the duplicates you cannot see

“Find records that are the same thing recorded twice — same customer, same date, same amount, different ID. Also flag names that are probably the same company spelled two ways.”

8. Explain the number to a human

“Summarise this data in five sentences for someone who will not open the file. Lead with the thing they most need to know, and flag anything you are unsure about.”

WHICH ONE WORKED? WHAT DID YOU CHANGE?

Practice prompts: meetings, service & planning

9. Meeting notes into action

“Turn these notes into three lists: decisions made, action items with an owner and a date, and open questions. If an owner or date is missing, say so instead of guessing.”

10. Answer the same question better every time

“Here are five ways customers ask about [topic] and how I usually answer. Write one clear answer I can reuse, plus two short variations for email and for the phone.”

11. Think it through before you commit

“I am considering [decision]. Give me the strongest case for, the strongest case against, and the three questions I should answer before deciding. Be direct.”

12. Learn something faster

“Explain [topic] to me the way you would to a smart person who has never worked in this industry. Then ask me three questions to check whether I actually understood it.”

WHICH ONE WORKED? WHAT DID YOU CHANGE?

Build your own prompt

Pick one real task. Fill in each line, then read the whole thing back as a single paragraph — that is your prompt.

The task, in one sentence

Role & goal — who is this helping, and what does “good” look like?

Context — what can the AI not possibly guess?

Task — the specifics, named out loud

Format — how do you want it back?

Guardrails — what should it not do?

What actually came back?

Build your own prompt

Pick one real task. Fill in each line, then read the whole thing back as a single paragraph — that is your prompt.

The task, in one sentence

Role & goal — who is this helping, and what does “good” look like?

Context — what can the AI not possibly guess?

Task — the specifics, named out loud

Format — how do you want it back?

Guardrails — what should it not do?

What actually came back?

Choosing a tool

Everything in this workbook works in all of them. Pick one and get good at it rather than sampling all five.

Tool	Free tier	Where it shines
Claude claude.ai	Yes	Long documents, careful reasoning, file analysis, saving a process as a reusable skill
ChatGPT chatgpt.com	Yes	The most widely used; strong general assistant with a large add-on ecosystem
Grok grok.com	Yes	Conversational, tied into X; a blunter tone some people prefer
Microsoft Copilot	Limited	Already inside Word, Excel, and Outlook if your company runs Microsoft 365
Google Gemini	Yes	Sits inside Gmail, Docs, and Sheets for Google Workspace shops

Honest advice: *the tool matters far less than the habits. A specific prompt in a free tool beats a lazy prompt in a paid one, every time.*

THE TOOL I AM COMMITTING TO FOR THE NEXT MONTH

Glossary

The words people use around AI, in plain English.

Prompt — what you type. The instruction and the context you give the AI.

LLM (large language model) — the technology under the hood. It predicts plausible text based on patterns, which is why it is good at language and shaky at arithmetic.

Context window — how much the AI can hold in mind at once. Very long conversations start to lose the early details.

Hallucination — a confident, well-written, entirely made-up answer. The reason you verify anything factual.

Grounding — giving the AI your actual file or data instead of letting it answer from general knowledge.

Token — a chunk of text, roughly three-quarters of a word. Usage limits are usually measured in tokens.

Memory — a setting that lets the tool carry details between chats. Useful at work, worth turning off during practice.

Glossary, continued

Project / Custom GPT — a saved workspace with your files and instructions already loaded, so you skip the setup every time.

Skill — a saved, repeatable procedure the AI can run on command. Where the process file you write in Step 4 eventually ends up.

Agent — an AI that takes multiple steps on its own, using tools, rather than answering one question at a time.

Human in the loop — a person reviews and approves before anything is final. The standard for anything touching money, people, or customers.

Fine-tuning — retraining a model on your own data. Expensive, rarely necessary, and usually suggested by someone selling it. Good prompts and good context solve most of what people reach for fine-tuning to fix.

Vibe coding — describing what you want in plain language and letting the AI build it, then steering by reacting to what comes back. What we are doing today.

WORDS YOU HEARD TODAY THAT YOU WANT TO LOOK UP

Your next 30 days

Nothing here takes more than twenty minutes. The point is repetition, not ambition.

- 1 This week** — run your Step 4 process file on one real task. Same loop, your data.
- 2 Week two** — pick three practice prompts from pages 23–25 and use them on live work. Note what you had to change.
- 3 Week three** — refine your process file with what you learned. Show it to one colleague and have them run it.
- 4 Week four** — write down your ground rules from page 22 and share them with your team. Decide what you would automate next.

One task I will run this loop on, and when

One person I will show this to

What would have to be true for this to be worth doing every week?

Thank you for coming

You now have a repeatable loop, a process file, and a stack of prompts to try on real work. The only thing left is the second run.

If you get stuck, or you want a hand pointing this at an actual workflow in your business — that is what we do. Reply to the follow-up email or reach out any time.

connect@kalflow.com

www.kalflow.com

Southwest Michigan

Innovation. Education. Community.