

AI SKILLS GUIDE

The AI Skills Guide

Become the 1% of AI Users

.TrackImpact



CONTENTS

In This Guide

1 — The Landscape	5
2 — Tool Choice	6
3 — Prompting Basics	8
4 — Your Context	10
5 — AI Judgment	11
6 — Real Work	13
7 — The Bridge	14

Become the 1% of AI Users

Almost everyone with access to AI chat has typed a question into it. A much smaller group has actually changed how they work because of it.

The gap between those two groups isn't access, and it isn't intelligence. It's a small set of habits — how you pick a tool, how you prompt, how you give AI context, how you check its work, and how you turn what it gives you into something usable.



None of these habits are complicated. Most of them take longer to read about than to actually do. But almost nobody does all of them, which is exactly why the gap exists and why it's closable.

This guide is those habits — six of them, in order, followed by one about applying all of it to your specific job.

Work through it once, then come back to whichever section you need next time you're stuck. It's built to be used that way, not just read once.

BEFORE YOU START

How to Use This Guide

Sections 1 through 6 are general — the same six habits whether you're in sales, finance, operations, or anywhere else. Section 7 is where it gets specific to your job.

Read it start to finish, or jump straight to whichever habit you're missing:

- 1 — Chat, Work, and Code — pick the right surface
- 2 — Tool Choice — a fast decision, not a memorised chart
- 3 — Talking to AI Well — guardrails, not tricks
- 4 — Giving AI Your Context — stop re-explaining yourself
- 5 — AI Judgment — trust, but verify
- 6 — Turning Output Into Real Work — the finishing step
- 7 — Your Specific Job — where the general becomes personal

The habits compound. The sooner you start using them — not just reading them — the sooner they pay off.

Chat, Work, and Code

AI tools aren't one thing. The same underlying model shows up in three very different products, and knowing which one you're in — and why — is the first habit.



Chat — Quick Help, No Setup

This is ChatGPT, Claude, Copilot Chat, and similar — open a window, ask a question, get an answer. No setup, no connection to your files or systems. Best for one-off questions, drafting, brainstorming, and anything where the AI doesn't need to know anything about your specific situation beyond what you type.

Work — AI That Knows Your Context

This is Claude Cowork, ChatGPT Work, Microsoft Copilot inside your Microsoft 365 apps, and similar — AI connected to your files, your email, your calendar, your company's systems. Slower to set up, far more useful once it's set up, because it isn't starting from zero every time you ask it something.

Code — Technical and Agentic Work

This is Claude Code, OpenAI Codex, and similar — AI that can write and run code, use tools, and carry out multi-step tasks with real actions attached to them, not just text back. Built for developers, but increasingly relevant to anyone automating a genuinely multi-step process.

You don't need to pick one and stick with it — most people who use AI well move between all three depending on the task. What matters is recognizing which one you're in, and not expecting chat-level tools to behave like work-level ones, or vice versa.

Choosing the Right Tool for the Job

Once you know the three categories, picking between them stops being a research project and becomes a handful of fast questions.



#1

Quick answer or first draft?

Use chat. No setup cost, and you'll have an answer before a work-tool would finish loading your context.

#2

Touches your files or company data?

Use a work tool. It already knows your context, so you're not re-explaining your situation in every prompt.

#3

Multiple steps, or taking an action?

Use a code or agentic tool. These can actually do things — run steps, call tools, carry a task through to completion — not just describe them.

#4

Need an image or video generated?

Use a dedicated generation tool. Chat and work tools can describe what you want; purpose-built tools are what actually make it.

None of this needs to be memorized as a chart. Ask yourself what the task actually requires — speed, context, action, or generation — and the right tool is usually obvious.

TRACKIMPACT — YOUR AI PLAYBOOK

By Job. By Task. By Your Actual Tool Stack.

TrackImpact's guidance is matched to three things at once: your job, the specific task in front of you, and which AI tools your company has actually licensed — so the answer isn't generic advice, it's what to do right now, with what you already have access to.

Every time that guidance gets used, it's logged automatically — building the record that shows leadership how AI is actually being used across the organization, and what impact it's having.

YOUR ROLE

Senior Financial Analyst

THE TASK

Reconcile month-end variance

YOUR COMPANY'S STACK

Excel Copilot, ChatGPT Enterprise

→ RECOMMENDATION

"Use Excel Copilot to flag variances outside 5%, then draft the exception note in ChatGPT Enterprise using your team's template."

LOGGED TO YOUR TRACKIMPACT RECORD

See how it works → www.trackimpact.io/

Talking to AI Well

If there's one skill that separates shallow AI use from real AI use, it's this one — and the good news is it doesn't take a framework.



You Don't Need a Prompting “Framework”

You'll see terms like one-shot, few-shot, chain-of-thought, or role prompting in AI tutorials — one-shot means giving a single example of what you want, few-shot means giving several, chain-of-thought means asking the AI to reason step by step before answering, and role prompting means asking it to respond as a specific kind of expert. They're useful vocabulary if you want it, but none of them is a prerequisite. AI chat has gotten meaningfully better at understanding plain language, so you don't need to pick a named technique before you type — you need to say clearly what you mean.

Set Guardrails With Three Things: Voice, Goal, Context

What actually keeps a response on track isn't a technique, it's giving the AI three pieces of information up front: the voice (who should it sound like — a formal analyst, a casual colleague, you), the goal (what is this response actually for — to inform, to persuade, to get approved), and the context (what does it need to know about the situation to avoid a generic answer). Leave any one of the three out and you'll usually get something technically correct but slightly off.

WEAK PROMPT

“Write an email about the delay.”

BETTER PROMPT

“Write a two-paragraph email to a client, apologetic but not over-explaining, telling them their shipment is delayed two weeks and giving the new date of March 15.”

The second version gives the AI its voice (apologetic, not over-explaining), its goal (inform without escalating), and its context (a client, a two-week delay, a specific new date). That's

the whole skill — the AI isn't reading your mind, so tell it what you'd tell a new hire doing the task for the first time.

Treat the First Response as a Draft, Not an Answer

The first response is a starting point, not a verdict. If it's close but not right, tell it exactly what to change — “shorter,” “more formal,” “cut the second paragraph” — rather than starting over with a new prompt. Iteration is faster than re-explaining from scratch, and it's how you actually get to a result you'd put your name on.

Let AI Ask You the Questions

One habit that feels strange the first time you try it: instead of writing the perfect prompt yourself, ask the AI to ask you clarifying questions before it answers — literally “ask me whatever you need to know before you start.” It will surface the exact ambiguities you didn't think to resolve: who's the audience, what's the constraint you haven't said out loud, what does “good” look like here.

It feels slower the first time. It isn't, really — and the questions it asks are usually the same handful of things you'll learn to just include in your first prompt next time. The back-and-forth is what teaches you to write a better first prompt going forward, not a permanent extra step.

Put together, prompting well is less a trick and more a habit: say the voice, goal, and context up front, treat the first answer as a draft, and let the AI ask when you're not sure what it needs. That's the whole skill — everything else is refinement.

Giving AI Your Context

The single biggest difference between a generic AI response and a genuinely useful one is whether the AI knows anything about your specific situation.



Most AI tools now have a way to load standing context so you're not re-explaining yourself every time: Claude has Projects and Skills, ChatGPT has Projects, custom GPTs, and memory, Microsoft Copilot grounds itself in your organization's Graph data — your files, emails, and calendar. Setting this up once pays off every single time you prompt after that.

The difference shows up immediately when you compare the same request with and without context loaded:

SAME PROMPT, NO CONTEXT

"Write a project update for my manager."

Generic structure, placeholder details, nothing your manager couldn't have written faster themselves.

SAME PROMPT, WITH CONTEXT LOADED

"Write a project update for my manager."

Pulls the actual project name, real milestones, real blockers, and matches the tone of your last few updates — because the AI already knows what project you're on.

Context isn't a one-time setup task, either — it's worth periodically checking what's loaded and pruning what's gone stale, the same way you'd clean out an old reference document.

One caveat: context makes AI more useful, not more correct. It still needs the same verification habits covered in Section 5 — loaded context reduces how often you need to explain yourself, it doesn't remove the need to check the output.

Trusting but Verifying

AI is confident by default — confident whether it's right or wrong. Knowing when to trust it and when to check it is the habit that keeps you out of trouble.



AI Can Be Confidently Wrong

A wrong answer from AI reads exactly like a right one — same fluent tone, same confident structure. There's no built-in signal for “I'm not sure about this part.” That's not a flaw to work around occasionally, it's a permanent property of how these tools work, and it's the reason verification is a habit, not an exception.

The Fix Is a Habit, Not a Tool

You don't need a special checking tool. You need a default: skim anything AI hands you for the parts that matter most before it goes anywhere — names, numbers, dates, anything you didn't personally supply as a fact in your prompt. If the AI generated a number rather than restating one you gave it, verify it before you use it.

Know Where AI Is Reliable Versus Where It Needs a Check

Where AI Actually Delivers

✓ **Drafting From Your Material**

✓ **Restructuring & Summarizing**

✓ **First-Pass Work**

Where It Still Needs a Human

✗ **Facts You Didn't Provide**

✗ **Anything With Real Consequences**

Get this habit right and everything downstream gets easier — including the tracking and evidence-building covered in Section 7, which depends on the work you're tracking actually being correct.

Turning AI Output Into Real Work

Getting a good response from AI is only half the job. Turning it into something you'd actually send, file, or ship is the finishing step most people skip.



Don't Copy-Paste Raw Output as the Deliverable

AI output is a draft, formatted the way a chat window formats things — not the way your deliverable needs to look. Move it into the right document, adjust it to your house style, and read it once as if you were the recipient before it goes anywhere.

Tell AI What Format You Actually Need, Upfront

Instead of getting generic prose and reformatting it yourself, ask for the structure you actually need — a table, a numbered list, a specific set of headers — in the original prompt. It's faster to specify than to restructure after the fact.

Know When to Let AI Structure Something Versus Writing Free-Form

Structured output (tables, lists, templates) is where AI is most reliable — the format itself constrains the answer. Free-form writing needs more of your own judgment layered on top, because there's nothing forcing it to stay on track.

Get comfortable with this finishing step and AI output stops being something you have to apologize for or quietly clean up — it becomes indistinguishable from work you did yourself, because by the time it's done, in the ways that matter, you did.

From General Skills to Your Specific Job

Everything so far applies no matter what you do for a living. This last section is about applying it to what you actually do for a living.



The six habits work on any task. But not every task in your job is worth automating, and not every automatable task is worth automating first — so before you start applying AI to your specific role, it helps to know how to pick where to start.

Before You Automate Anything, Run This Checklist

Run any recurring task through these four questions before you decide it's worth handing to AI:

✓ **What does an hour of this task actually cost?**

Multiply the time it eats by the loaded cost of the person doing it. Twenty minutes a day from a junior analyst is a different priority than the same twenty minutes from someone billed out at three times the rate — cost, not just time, is what should rank your list.

✓ **How many hours does it eat across the whole team, not just you?**

One person saving fifteen minutes a day is nice. Ten people on the same recurring task saving fifteen minutes each is a meaningful chunk of an FTE handed back — freed time that can go toward work that actually generates revenue, not just admin.

✓ **Is this a Human-in-the-Loop (HITL) task, or can it run on its own?**

HITL means AI does the work but a person still reviews it before it counts — the right call for anything involving judgment, nuance, or real consequences if it's wrong. Repetitive, rules-based, low-stakes tasks can run with far less oversight. Knowing which bucket a task falls into tells you how much automation is actually realistic, not just possible.

✓ **Which KPI would actually move if this got faster or better?**

Automating for its own sake doesn't prove anything. Tie the task to a number your role is already measured on — response time, error rate, close rate, throughput — so that when it's done, there's a real answer to whether it helped.

That last question is the one TrackImpact's job-profile search is built to answer directly — instead of guessing your way through the checklist above, you can look up your actual job title and see the real, task-by-task breakdown already mapped out.

TRACKIMPACT — SEARCH PROFILES

Not Sure Where to Start? Search Your Job Title.

TrackImpact's marketplace covers thousands of O*NET job titles, and every profile breaks the role down into its real, recurring tasks — not "use AI for marketing," but the actual activities someone in that seat does every week — each matched to a concrete AI approach and an automation-readiness read.

Instead of guessing which tasks are worth automating, search the job title and see the answers already mapped out.

SEARCH

"Senior Financial Analyst"

TASKS MAPPED

38 recurring tasks found

→ SAMPLE TASK MATCH

"Reconcile the variance report against last month, flag anything outside a 5% band."

HITL SCORE: MEDIUM — AI DRAFTS, YOU REVIEW

Search your job → www.trackimpact.io/search/

Proof That Pays Off Twice

For the organization, this is what makes AI adoption measurable instead of anecdotal — leadership can see which teams and workflows are actually generating value from AI, not just which ones say they are, and make rollout and budget decisions based on real usage instead of enthusiasm.

For you individually, it's career-proofing: as AI reshapes how work gets done, the people who come out ahead won't just be the ones using it — they'll be the ones who can prove they used it well. A personalized report turns your own AI use into a track record you can point to in a review or a promotion conversation. Team reports do the same thing one level up, showing where a whole group is creating value and where there's still room to close the gap.

To be precise about what's here today versus what to expect: the job-profile matching, activity tracking, and personalized and team reporting described here are live, real features, not a roadmap pitch.

TRACKIMPACT — TRACK & SCORE

Turn Everyday AI Use Into Evidence

Once a task is worth doing, TrackImpact logs it automatically — the tool, the task, the time saved. From there it's built to answer three questions at a glance: where you're already using AI well, what to do next, and how you actually benefited. Every entry is scored for automation-readiness, and the highest scorers are flagged as candidates to become a standing loop.

Outcomes are the report's whole point — not activity for its own sake. Track, score, and step back is the mechanism; what it hands back is a clear, evidence-backed answer to “did this actually help,” every time.

TOOL	TASK	TIME SAVED	SCORE
ChatGPT	Draft first-pass client email	15 min	82 — Loop candidate
Otter.ai	Transcribe client call	20 min	65
Excel Copilot	Reconcile month-end entries	35 min	40



↳ Recorded outcome feeds the next run


See your Track & Score → dev.trackimpact.io/dashboard

TRACKIMPACT — TEAMS

See Impact Across the Whole Organization

Teams brings every individual's tracked and scored activity together at the organizational level — so leadership can see adoption and value by department, site, or role, and compare who's driving results.

It's the difference between a handful of enthusiastic early adopters and a rollout that can actually be managed and reported on.



See how it works → trackimpact.io

Get Started

Find your role: trackimpact.io/search

Start tracking: [your TrackImpact dashboard](#) — [Track Impact](#)



Ade Molajo

Founder, TrackImpact

LinkedIn: linkedin.com/in/ademolajo

Email: ade.molajo@trackimpact.io

Web: trackimpact.io

READ NEXT

The AI Impact ROI Guide

This guide built the individual habits. The AI Impact ROI Guide is the organizational half — a maturity model, self-assessment, and ROI framing for scaling AI adoption across sales, finance, marketing, and people functions.

Built for AI transformation leads and Chief AI Officers who need to prove impact company-wide, not just build good personal habits.



Read the ROI Guide → www.trackimpact.io/guides/ai-impact-roi/

Photo Credits

All photography sourced from Unsplash. The Unsplash License does not require attribution, but it's listed here as good practice.

- Team meeting in a modern office — Vitaly Gariev / Unsplash (img-skills-opening.jpg)
- Person working across a laptop and notes — Unsplash (img-skills2-landscape.jpg)
- Close-up of hands typing on a laptop keyboard — Unsplash (img-skills2-toolchoice.jpg)
- Person typing on a laptop during a focused work session — Unsplash (img-skills2-prompting.jpg)
- Colleague explaining background and context at a desk — Vitaly Gariev / Unsplash (img-skills-context.jpg)
- Two colleagues reviewing documents together — Vitaly Gariev / Unsplash (img-skills-judgment.jpg)
- Team collaborating around a whiteboard during a meeting — Vitaly Gariev / Unsplash (img-skills-output.jpg)
- Person at a desk with laptop and papers — Unsplash (img-skills-bridge.jpg)
- The AI Impact ROI Guide cover — Vitaly Gariev / Unsplash (img-generic-cover.jpg)
- Person thinking closely at a laptop by a bright window — Unsplash (img-skills2-cover.jpg)