POST-SESSION RESOURCE

Practical AI for Real Work

Your companion guide to the fundamentals, prompting, workflows, and tools — built to turn one session into a lasting working advantage.

🎯 **Who it's for:** Beginners to intermediate professionals putting AI to work in real jobs.

⏱ **Read time:** ~20 minutes · keep it open as a reference, not a one-time read.

🧰 **Inside:** 14 key concepts, a prompting playbook, 8 workflows, a 14-tool toolkit, and a 7-day plan.

🧭 **The goal:** Cut through the noise, stay ahead, and out-skill the room.

SECTION 01

Session overview

This session was about making AI genuinely useful in your day-to-day work — not as a novelty, but as a dependable tool you reach for the way you reach for a spreadsheet or a search bar.

The journey ran from the ground up: first **how these tools actually work** under the hood (in plain language), then **how to talk to them well** through better prompting, then **how to choose the right tool** for a given job, and finally **how to wire AI into real workflows** — meetings, research, analysis, content, and more.

Led by **Raghuveer (RV)** — a marketer, product person, and AI generalist — the framing was deliberately practical. You don't need to be an engineer to get disproportionate value from AI. You need a clear mental model, a few reliable habits, and a toolkit you actually use.

> **Why this matters**
>
> AI capability is now a skill gap, not just a tooling gap. The people who pull ahead aren't the ones with secret tools — they're the ones who've built the habit of delegating the right work to AI, checking it well, and moving faster as a result. This guide is built to help you become one of them.

SECTION 02

Key concepts, explained simply

You don't need the math. But a working mental model of what's happening makes you dramatically better at using these tools — and at knowing when to trust them.

Generative AI

AI that **creates** new content — text, images, audio, code — rather than just sorting or labelling existing data. Think of it as the difference between a librarian who finds you a book and a writer who drafts you a new one on request.

LLMs (Large Language Models)

A large language model is a **subset of generative AI focused on language**. It has read an enormous amount of text and learned the patterns of how language fits together, so it can read what you write and produce a fluent, relevant response. ChatGPT, Claude, and Gemini are all LLMs at their core.

Tokenization

Before a model can work with your words, it breaks them into **tokens** — small chunks, roughly a word or part of a word.

> *Analogy:* Think of building with LEGO. You don't think in finished houses — you think in bricks. Tokens are the model's bricks. "Workshop" might be one brick, or two ("work" + "shop"). The model only ever sees and arranges these bricks.

Embeddings

Each token gets converted into a list of numbers — an **embedding** — that captures its meaning. Words used in similar ways end up with similar numbers, so the model understands that "king" and "queen" are related, or that "invoice" and "receipt" live in the same neighbourhood.

> *Analogy:* Imagine a giant map of meaning. Every word is a pin on that map. "Coffee" and "espresso" sit close together; "coffee" and "tractor" sit far apart. Embeddings are the coordinates of each pin.

Attention (the transformer mechanism)

This is the breakthrough that made modern AI work. **Attention** lets the model weigh which earlier words matter most when interpreting the current one.

> *Analogy:* In the sentence "the trophy didn't fit in the suitcase because **it** was too big," you instantly know "it" means the trophy. Attention is how the model does the same thing — it learns where to look.

Prediction

At its heart, an LLM does one thing extraordinarily well: **predict the next token**. Given everything so far, it asks "what comes next?" and picks the most fitting continuation — then repeats, one token at a time.

Response generation

Stringing those predictions together — token after token — produces the full reply you read. It feels like the model is "thinking out a paragraph," but it's actually composing it one brick at a time, each choice informed by everything before it.

> **The one-line version**
>
> Your words become **tokens** → tokens become meaning-rich **embeddings** → **attention** decides what matters → the model **predicts** the next token → repeated prediction becomes a full **response**.

SECTION 03

What you learned

If you take nothing else away, take these:

- AI is a **spectrum of skill**, not a single button — there's a clear ladder from casual user to builder, and you can climb it deliberately.
- LLMs work by **tokenizing, embedding, attending, and predicting** — knowing this makes you a sharper, more realistic user.
- **Prompting is a craft.** A repeatable structure (role, context, task, format, constraints) beats a vague one-liner every time.
- Different assistants have **different sweet spots** — choosing the right one for the job is itself a skill.
- **Reasoning models** change what's possible for multi-step, high-stakes work — and they pair beautifully with written SOPs.
- The biggest wins come from **workflows, not one-off chats**: meetings, research, analysis, content, and your own files.
- A **focused toolkit** you actually use beats a bookmark folder of 50 tools you don't.
- Judgement, verification, and privacy awareness are **part of the skill**, not an afterthought.

SECTION 04

The AI skill ladder

AI ability isn't binary. It's a ladder, and most professionals can climb two or three rungs faster than they expect. Find where you are today — then look at the "how to move up" column.

From casual use to deep expertise

Level	What it means	What they typically do	How to move up
Basic users	Use AI occasionally for simple questions and quick drafts.	Ask one-off questions; copy-paste answers; treat it like a smarter search box.	Adopt a prompt structure; start giving context and a desired format.
Prompters	Get reliably good output through deliberate prompting.	Use roles, examples, and constraints; iterate on weak answers; save prompts that work.	Turn winning prompts into reusable templates and Projects; chain steps together.
Automations	Remove repetitive work by connecting steps and tools.	Use built-in connectors, scheduled tasks, and simple no-code automations.	Map a recurring process end-to-end; hand the routine parts to AI on a trigger.
Agents	Let AI take multi-step actions toward a goal, with tools.	Use agentic features that browse, read files, or operate apps to complete a task.	Define clear goals, guardrails, and check-points; learn to supervise, not micromanage.
Builders	Create their own AI-powered tools and apps.	Combine APIs, prompts, and logic into small products or internal tools.	Learn the basics of APIs and a little code (or vibe-code with AI's help).
AI generalists	Fluent across the whole landscape; match the right approach to any problem.	Move easily between tools, models, and workflows; advise others; spot what's worth doing.	Go deep on a domain (e.g. data, marketing ops) and ship repeatable systems.
ML engineers	Build, fine-tune, and deploy models in production.	Work with training pipelines, evaluation, infrastructure, and scaling.	Formal study of ML engineering, MLOps, and systems design.

Level	What it means	What they typically do	How to move up
Researchers	Push the frontier — invent new methods and architectures.	Publish, experiment, and advance what models can fundamentally do.	Deep academic and research specialisation over years.

> **Where to aim**
>
> For most working professionals, the sweet spot is climbing confidently to **Prompter** → **Automations** → **Agents**. That's where the time savings become dramatic without needing to become an engineer. This guide is built to get you there.

SECTION 05

The prompting playbook

A weak prompt gets a generic answer. A structured prompt gets something you can actually use. Here's a reusable framework inspired by the session's "Magic Prompt" idea — fill in the blocks and you'll rarely write a bad prompt again.

The Magic Prompt template

```
[ROLE]       Act as a <specific role / expertise>.
[CONTEXT]    Here's the situation and background: <who, what, why, audience>.
[TASK]       I need you to <the precise thing you want done>.
[FORMAT]     Give it to me as <email / table / bullets / 200 words / etc.>.
[CONSTRAINTS]Keep in mind: <tone, length, must-haves, must-avoids>.
[EXAMPLES]   Here's an example of the style/quality I want: <sample>.
```

You won't always need all six blocks — but **role, context, task, and format** are the everyday workhorses. Add constraints and examples when quality really matters.

The session example, done well

The session's worked example: writing an email to 10,000 subscribers inviting them to a Generative AI workshop. Watch the difference structure makes.

> ✘ **Weak prompt:** *"Write an email inviting people to my AI workshop."*
>
> ✔ **Structured prompt:** *"Act as an email marketer for a practical-AI brand. Context: I'm inviting 10,000 existing subscribers — busy professionals, mixed AI experience — to a free 90-minute Generative AI workshop next Thursday. Task: write the invitation email. Format: subject line + preview text + ~150-word body + one clear call-to-action button label. Constraints: warm but not hypey, emphasise practical takeaways, one CTA only, skimmable. Avoid jargon and exclamation overload."*

5 examples for real work

- **Inbox triage:** *"Act as my executive assistant. Here are 6 emails (pasted below). Summarise each in one line, flag anything urgent, and draft a one-paragraph reply to each. Format: a table with columns Sender · Summary · Urgency · Draft reply."*
- **Meeting prep:** *"Act as a strategy advisor. Context: I'm meeting a potential partner tomorrow (details below). Task: give me 5 smart questions to ask, 3 risks to probe, and 2 things to avoid saying. Format: short bullets."*
- **Document rewrite:** *"Act as a clear-writing editor. Rewrite the text below for a non-technical audience. Keep it under 200 words, use plain language, and preserve all numbers exactly."*
- **Data sense-making:** *"Act as a data analyst. Here's a table of last quarter's sales. Task: tell me the 3 most important trends, anything surprising, and what to investigate next. Format: bullets, no jargon."*
- **Learning fast:** *"Act as a patient tutor. Explain <topic> to me at three levels: a one-sentence version, a paragraph for a beginner, and the key nuance an expert would add."*

Common mistakes

- **Too vague.** "Make it better" gives the model nothing to aim at.
- **No role or audience.** The model defaults to bland, average output.
- **No format specified.** You get a wall of text when you wanted a table.
- **Everything in one giant prompt.** Break complex tasks into steps.
- **Giving up after one try.** The second prompt — "good, but shorter and warmer" — is where the magic happens.
- **Not showing an example.** One sample of the style you want is worth a paragraph of description.

How to improve a weak prompt

✘ **Before:** *"Summarise this report."*

✔ **After:** *"Summarise this report for a busy executive who has 60 seconds. Give me: the single key takeaway, 3 supporting points, and 1 recommended action. Plain language, under 120 words."*

Make it reusable: Claude Projects

When you find yourself prompting for the same kind of task repeatedly — weekly reports, customer replies, content in your brand voice — set up a **reusable Project** in Claude. Load it once with your context, instructions, and reference files, and every new chat in that Project already knows the rules. It turns a great one-off prompt into a permanent capability.

SECTION 06

Choosing the right AI tool

The major assistants overlap a lot — any of them will handle everyday tasks well. The differences below are practical leanings, not hard rules. The best move is to try two or three on your real work and notice which fits your style.

The major assistants at a glance

Tool	Best use cases	Strengths	When to try it
ChatGPT	General-purpose tasks, brainstorming, drafting, everyday Q&A.	Versatile all-rounder with a broad feature set and large ecosystem.	A reliable default for almost anything.
Claude	Long-form writing, careful reasoning, working with documents and files.	Thoughtful, natural writing; strong on nuance; Projects for reusable context.	When quality of writing and reasoning matters most.
Gemini	Tasks tied to Google's ecosystem and multimodal inputs.	Integration with Google Workspace; handles text, images, and more.	If you live in Gmail, Docs, and Drive.
Copilot	Work inside Microsoft 365 apps.	Embedded in Word, Excel, Outlook, and Teams.	If your workplace runs on Microsoft.
Meta AI	Casual, conversational use across Meta's apps.	Built into WhatsApp, Instagram, and Messenger.	For quick help where you already chat.
Grok	Conversational answers with a real-time, social-media flavour.	Tied closely to X; informal tone.	For current, of-the-moment topics.
Perplexity	Research questions where you want sources.	Answers come with citations you can verify.	When you need to trust and trace the answer.

A simple rule of thumb

Reach for a **research tool with citations** when you need to verify facts, a **writing-and-reasoning assistant** when quality of thought matters, and **whatever's embedded in your daily software** when convenience wins. Don't overthink it — fluency comes from using them, not from picking the "perfect" one.

SECTION 07

Reasoning models & SOPs

What reasoning models are

Most LLMs answer fast — they predict their way straight to a reply. **Reasoning models** are designed to slow down and work through a problem in steps before answering, almost like showing their working. That extra "thinking time" makes them noticeably better at multi-step, logical, and high-stakes tasks.

Why they matter

- They handle **complex, multi-step problems** — analysis, planning, debugging, math-like logic — far more reliably.
- They make **fewer careless errors** on tasks where one wrong step ruins the answer.
- They're better at **following detailed instructions** and long SOPs faithfully.

When to use them

Use a reasoning model when the task is **important, layered, or easy to get subtly wrong**: planning a project, analysing data, reviewing a contract's logic, or building a process. For quick drafts, simple rewrites, or brainstorming, a faster standard model is fine — and snappier.

What an SOP for AI looks like

An SOP (Standard Operating Procedure) is just a **written recipe** for a recurring task. When you combine a clear SOP with a capable model, you get consistent, repeatable output every time — instead of re-explaining yourself in every chat. Tools like Claude and coding assistants such as **Codex** can follow these procedures, and structured `skill.md` files let you package an SOP so a model picks it up automatically.

SOP template 1 — Weekly report

```
GOAL: Produce my Monday status report.
INPUTS: Last week's notes + this week's task list (pasted/attached).
STEPS:
  1. Summarise what shipped last week in 4 bullets.
  2. List this week's top 3 priorities.
  3. Flag any blockers or risks.
OUTPUT: One email, under 200 words, plain professional tone.
```

SOP template 2 — Customer reply

```
GOAL: Draft a reply to an inbound customer message.
INPUTS: The customer's message + relevant order/account details.
STEPS:
  1. Identify the customer's core question and emotional tone.
  2. Answer clearly; if unresolved, state the next step + timeline.
  3. Match our brand voice: warm, concise, no jargon.
OUTPUT: Ready-to-send reply + a one-line internal note on what I changed.
```

SOP template 3 — Research brief

```
GOAL: A decision-ready brief on a topic.
INPUTS: The question + any links or constraints.
STEPS:
  1. Gather key facts from credible sources (with citations).
  2. Note points of disagreement or uncertainty.
  3. Give a recommendation with the reasoning behind it.
OUTPUT: 1 page — Summary, Evidence, Open questions, Recommendation.
```

Example `skill.md`-style structure

A `skill.md` packages a reusable capability so an AI tool knows *when* and *how* to apply it:

```
# Skill: Brand-Voice Email Writer

## Description
Use this skill whenever I ask to draft or rewrite a marketing
email. Triggers: "write an email", "draft a newsletter".

## Voice & rules
- Warm, clear, confident. Never hypey. One CTA only.
- Short paragraphs. Skimmable. No exclamation pile-ups.

## Process
1. Confirm audience, goal, and the single CTA.
2. Draft subject + preview text + body (~150 words).
3. Offer one alternative subject line.

## Output format
Subject · Preview · Body · CTA label
```

SECTION 08

Practical AI workflows

This is where the real value lives. Each mini-guide below turns a session demo into a repeatable workflow you can run this week.

1. Meeting notes / AI meeting assistant

- **When to use it:** Any call you'd otherwise scramble to take notes in, or can't attend live.
- **Suggested tools:** Fireflies for capture; Claude or ChatGPT to refine the output.

1. Connect the notetaker to your calendar so it joins meetings automatically.
2. Let it record, transcribe, and produce a summary with action items.
3. Paste the transcript into an assistant for a cleaner recap or follow-up email.

Example prompt: "From this meeting transcript, give me: a 3-line summary, a table of action items (owner · task · due date), and a short follow-up email to attendees."

2. Interview prep with voice mode

- **When to use it:** Rehearsing for an interview, a tough conversation, or a high-stakes pitch.
- **Suggested tools:** Voice mode in ChatGPT or Gemini.

1. Tell the AI the role, company, and the kind of conversation you're preparing for.
2. Ask it to interview you out loud, one question at a time.
3. Request feedback after each answer — clarity, confidence, what to cut.

Example prompt: "Act as a hiring manager for a product role. Interview me by voice, one question at a time. After each answer, give me 20 seconds of honest feedback before the next question."

3. Deep research

- **When to use it:** You need a thorough, sourced answer — not a quick guess.
- **Suggested tools:** Perplexity, or the deep-research modes in Claude, ChatGPT, or Gemini.

1. State the question precisely, including scope and what a good answer contains.
2. Let the tool gather and cite sources; skim the citations for credibility.
3. Ask follow-ups to fill gaps, then request a structured summary.

Example prompt: "Research <topic> for a business decision. Give me the current state, the main options with trade-offs, and a recommendation — with citations I can check."

4. Dashboard simplification

- **When to use it:** A dashboard or report is dense and you need the story fast.
- **Suggested tools:** Claude or ChatGPT (share a screenshot or export).

1. Upload a screenshot or the underlying export of the dashboard.
2. Ask what the key numbers mean and what's changed.
3. Ask for the one or two things that deserve action.

Example prompt: "Here's a screenshot of our analytics dashboard. Explain what it's telling me in plain English, highlight anything unusual, and tell me the two things I should act on."

5. Data analysis with AI

- **When to use it:** You have a spreadsheet and questions, but not the time to pivot-table it.
- **Suggested tools:** ChatGPT or Claude for analysis; Numerous.ai to work inside the sheet itself.

1. Upload the file (or describe the columns) and state your question.
2. Ask for trends, outliers, and a plain-language interpretation.
3. Request a chart or a summary table you can drop into a report.

Example prompt: "Here's our sales data. Find the top 3 trends, flag anything surprising, and suggest what to investigate next. No jargon."

6. Content / email drafting

- **When to use it:** Recurring writing — newsletters, posts, outreach, announcements.
- **Suggested tools:** Claude Projects for brand voice; Supergrow for LinkedIn content.

1. Give the assistant your audience, goal, and voice (an example helps).
2. Draft, then refine with specific feedback ("shorter," "warmer," "stronger hook").
3. Save the winning setup as a Project so next time starts from your standard.

Example prompt: "Using the brand voice in this Project, write a LinkedIn post about <topic>. Hook in the first line, 3 short insights, one takeaway. No hashtags-as-clutter."

7. Working with local files

- **When to use it:** The answer lives in your own documents, PDFs, or notes.
- **Suggested tools:** Claude (with Connectors / file uploads); NotebookLM for grounded research across your sources.

1. Upload the files or connect the source you want the AI to read.
2. Ask questions answered strictly from those documents.
3. Ask for a summary, comparison, or audio overview of the material.

Example prompt: "Based only on the documents I've uploaded, answer: what are the key obligations in this contract, and where exactly does each one appear?"

8. Browser assistant workflows

- **When to use it:** You're reading on the web and want to act on the page without copy-pasting.
- **Suggested tools:** Claude Extension on Chrome.

1. Open the page you're working with — an article, a long doc, a dashboard.
2. Open the assistant in the sidebar and ask it to read and summarise.
3. Ask it to extract, draft, or take the next step based on what's on screen.

Example prompt: "Summarise this page in 5 bullets, pull out any dates and figures into a table, and draft a 3-line note I can send my team about it."

SECTION 09

The AI toolkit

Fourteen tools worth knowing, each with a clear job. Don't adopt all of them — pick the two or three that fit your work and go deep. Every name links to its official site.

Your post-session toolkit

Tool	Category	Best for	Practical use-case	Try this starter
Wispr Flow	Voice → text	Speaking instead of typing, anywhere on your computer.	Dictate emails, prompts, and notes far faster than you type.	*Dictate a messy first draft, then ask an AI to clean it up.*
Fireflies	Meeting notes	Capturing and summarising meetings automatically.	Auto-join calls, transcribe, and produce action items.	*Let it join your next call; review the summary after.*
NotebookLM	Research assistant	Answering questions grounded in your own sources.	Upload reports/PDFs and query them; generate audio overviews.	*Upload 3 documents and ask for a combined summary.*
Phot AI	Image / creative	Product and photo editing without a designer.	Remove backgrounds, enhance shots, create product imagery.	*Turn a plain product photo into a clean studio shot.*
ChatGPT	General assistant	An everyday all-rounder for most tasks.	Draft, brainstorm, analyse, explain, and prep.	*Ask it to be your assistant for one whole workday.*
Claude	General assistant	Quality writing, reasoning, and document work.	Long-form drafting; Projects for reusable context; file analysis.	*Build a Project for a task you repeat weekly.*
Gemini	General assistant	Working within Google's ecosystem; multimodal input.	Help across Workspace; reason over text and images.	*Ask it to summarise a long doc from your Drive.*

Tool	Category	Best for	Practical use-case	Try this starter
Perplexity	Answer engine	Research where you want sources you can verify.	Fact-find with citations; compare options credibly.	*Ask a research question and check every citation.*
Supergrow	LinkedIn content	Personal branding and consistent posting.	Generate posts and carousels; repurpose long content.	*Turn a blog post into a week of LinkedIn posts.*
Get Multi	Model comparison	Seeing several models answer the same prompt.	Compare outputs side-by-side to pick the best.	*Run one important prompt across models and compare.*
Happenstance	Network search	Finding the right person in your network.	Search your connections for warm intros.	*Ask who in your network could intro you to <company>.*
Numerous.ai	Spreadsheets	Bringing AI into Sheets and Excel via formulas.	Categorise, extract, and write at scale across rows.	*Auto-categorise a column of customer feedback.*
There's An AI For That	Tool directory	Finding a tool for a specific task.	Discover purpose-built AI tools by use-case.	*Search for a tool for a task you do often.*
Claude for Chrome	Browser assistant	AI help on whatever page you're viewing.	Read, summarise, and act on web pages in a sidebar.	*Summarise a long article without leaving the page.*

Recommended starter stacks

Start small. Here are four focused stacks depending on what you do most.

> 🌱 **For beginners:** ChatGPT or Claude as your daily assistant, plus There's An AI For That to discover tools as needs arise. One assistant, used daily, beats ten you dabble with.
>
> ✍️ **For content creators:** Claude (voice & long-form) + Supergrow (LinkedIn) + Phot AI (visuals), with Wispr Flow to draft by voice.
>
> 🔬 **For researchers:** Perplexity (sourced answers) + NotebookLM (your own documents) + a reasoning model in Claude or ChatGPT for synthesis.
>
> 📊 **For operators & analysts:** Numerous.ai (in-spreadsheet) + Fireflies (meetings) + Claude or ChatGPT for analysis and reporting.

SECTION 10

Your 7-day action plan

Skill comes from reps, not reading. Spend 20–40 minutes a day for one week and you'll feel the difference. Use real work, not toy examples.

One week to a working habit

Day	Focus	Do this (20–40 min)
Day 1	Pick your assistant	Choose one assistant and use it for every small question all day. Notice where it helps.
Day 2	Master the Magic Prompt	Take 3 real tasks and write each as a structured prompt (role · context · task · format).
Day 3	Improve weak prompts	Take yesterday's outputs and refine them with follow-ups. Save the two best prompts.
Day 4	Run a workflow	Pick one workflow from Section 08 (meetings, research, or data) and run it end-to-end.
Day 5	Build something reusable	Create a Claude Project or a short SOP for a task you repeat weekly.
Day 6	Add a specialist tool	Adopt one toolkit tool that fits your role; complete one real task with it.
Day 7	Reflect & plan	Note your 3 biggest time-savers and the one habit you'll keep. Schedule it into next week.

SECTION 11

Best practices & cautions

Using AI well means using it **responsibly**. These habits separate confident users from careless ones.

Verify before you trust

> ⚠️ AI can sound completely confident while being wrong — sometimes inventing facts, figures, or citations. For anything that matters, **verify the claim** against a primary source. Tools with citations (like Perplexity) make this easier, but the responsibility is always yours.

Protect privacy & sensitive data

> ⚠️ Don't paste confidential, personal, or regulated information into consumer AI tools unless you know how that data is handled. When in doubt, anonymise it, summarise it, or check your organisation's policy first.

Know when not to rely on AI

> ⚠️ Be cautious using AI as the final word on legal, medical, financial, or safety-critical decisions, or anywhere a subtle error carries real consequences. Use it to draft, explore, and accelerate — then bring in a qualified human to decide.

Keep human judgement in the loop

AI is a powerful assistant, not a replacement for your thinking. The best results come from a simple loop: **you direct, AI drafts, you review and decide.** Your taste, context, and accountability are exactly what the model lacks — and exactly what makes the output yours.

SECTION 12

Recommended next steps

- **Pick one assistant and go deep** for the next month before adding others. Depth beats breadth early on.
- **Build a personal prompt library** — a simple note of the prompts that consistently work for you.
- **Turn your best prompts into Projects or SOPs** so good results become repeatable, not lucky.
- **Automate one recurring task** this month — a weekly report, a content draft, a research routine.
- **Teach someone else** what you learned. Explaining it is the fastest way to lock it in.
- **Stay lightly current.** Follow one or two trustworthy sources rather than chasing every announcement.

SECTION 13

Quick-reference cheatsheet

Everything on one screen. Keep this handy.

🧠 Key concepts

Tokens = the bricks · Embeddings = meaning as coordinates · Attention = what matters most · Prediction = next-token guessing · Response = many predictions, strung together.

✍️ Prompt formula

Role → **Context** → **Task** → **Format** → **Constraints** → **Example.** The everyday four are **Role, Context, Task, Format.**

🎯 Best use cases by tool

Quality writing & reasoning → **Claude.** Everyday all-rounder → **ChatGPT.** Google ecosystem → **Gemini.** Microsoft 365 → **Copilot.** Sourced research → **Perplexity.** Your own documents → **NotebookLM.**

> **⚙️ Workflow ideas**
>
> Meetings → notetaker + recap · Research → cited deep-dive · Data → upload & ask · Content → Project + brand voice · Files → ground answers in your docs · Web → browser sidebar.

> **🧰 Tool suggestions**
>
> Start with one assistant (**Claude** or **ChatGPT**). Add **Perplexity** for research, **Fireflies** for meetings, **Numerous.ai** for spreadsheets, and **Wispr Flow** to work by voice.

> **⚠️ Always remember**
>
> Verify what matters · guard sensitive data · keep a human in the loop on big decisions.

A final word.

There's a lot of noise in AI right now — endless tools, endless hype, endless takes. The people who win aren't the ones who chase all of it. They're the ones who learn the fundamentals once, build a few reliable habits, and quietly compound the advantage week after week.

You now have the map: how these tools think, how to talk to them, which to reach for, and how to fold them into real work. The only thing left is reps.

> **Cut through the noise. Stay ahead. Out-skill the room.** Not by knowing every tool — but by being the person who actually puts AI to work, every single day. Start with Day 1. The rest follows.

Thank you for being part of the session. Now go build the habit.