

Putting AI to Work

What to do Monday morning

Alicia Dahling, MBA · Founder & CEO, Dahling Consulting Prepared for The Fleck Forum · Wayan Aspen · August 12, 2026

Thank you for an hour of genuine argument. This is the short version — five pages you will actually read. The long version is eighteen pages and covers every major tool with current pricing, how to keep your data out of model training, how to run these models on your own laptop so nothing touches a cloud service, and a section for each of the roles that were in the room today. It is free at **DahlingConsulting.com/ai**.

Everything below is something you can do this week without hiring anyone.

The three questions

If you keep nothing else from today, keep these. They resolve almost every practical decision about when to use these tools and when to put them down.

One. Am I using this to prepare for a conversation, or to avoid one?

You will know the answer before you finish asking it. Preparing is what these tools are extraordinary at — rehearse the hard conversation, draft the difficult email, stress-test the argument you are about to make. Avoiding is the failure mode, and it is quiet. Nobody decides to stop talking to their colleagues. The path of least resistance simply moves.

Two. Did the effort matter here?

Usually it did not. Nobody was ever moved by the effort in a meeting agenda or a board packet summary. Automate those without hesitation. But for a condolence note, an apology, or a difficult piece of praise, the effort is the message. A machine can produce better sentences in four seconds. What it cannot do is give away your twenty minutes, and the twenty minutes was the gift.

Three. When did this thing last tell me I was wrong?

If the answer is never, you are not using a tool. You are keeping company with a flatterer. Research published in Science in March of this year, across eleven major models, found that AI affirms a user's position markedly more often than another person would — and that exposure made people less willing to repair conflicts and more convinced they were right. Every good thing any of us has made was made better by somebody telling us it was not good yet.

Which tool for which job

You do not need all of these. You need one paid subscription and the knowledge of when to reach for something else. Prices are as of this month and will change.

If you need to	Use	Cost
Think through something hard, write, analyze	Claude or ChatGPT — pick one and learn it well	~\$20/month
Research a question with real citations you can check	Perplexity	Free tier is genuinely usable; ~\$20/month for more
Work inside documents and email you already own	Microsoft Copilot or Google Gemini	Bundled or ~\$20-30/month
Build a slide deck from your own notes fast	Gamma	Free tier available
Keep everything on your own machine, nothing in the cloud	Ollama or LM Studio (see below)	Free
Transcribe and summarize meetings	Built-in tools in Teams, Zoom, or Google Meet	Usually included

The honest recommendation. Start with one paid subscription — Claude or ChatGPT — and use it daily for two weeks before adding anything. The most common expensive mistake I see is buying four tools and learning none of them.

One caveat on the bundled assistants. Copilot and Gemini are convenient because they sit inside files you already own, but they are generally weaker than the standalone tools at hard reasoning and long-document work, and help with complex multi-sheet spreadsheets degrades quickly.

How to ask for something and get it

Most disappointing results come from a request that would have failed if given to a new employee. The pattern below fixes the majority of them.

Role. "You are reviewing this as a skeptical buyer." **Task.** What you want, specifically. Not "help with my listing" but "write three versions of this listing description, each under 120 words." **Context.** Who it is for, what came before, what must not appear. **Format.** Length, structure, tone. Ask for a table if you want a table. **Constraint.** What to avoid. "No marketing language. No exclamation points. Do not compliment me."

Then do the thing almost nobody does: argue with it. The first answer is a draft, not a product. Three follow-ups that consistently improve output:

"What is the strongest objection to this?"

"What did you assume that I did not tell you?"

"Rewrite this as if you disagreed with me."

Set your standing instructions once. Every major tool has a settings field for permanent instructions. Mine says: I value directness and brevity. Do not open with compliments or restate my question. When you are uncertain, say so plainly rather than guessing. When I ask for analysis, give me the strongest objection to my position as well as the support for it. Use plain American business English, never marketing language. This single step improves every future answer.

If you would rather it not be in the cloud

A number of you raised this, and it is a legitimate position rather than a paranoid one. Two things are worth separating: whether a company trains its models on what you type, and whether your data sits on someone else's server at all.

Turn off training in what you already use. In ChatGPT this lives under Settings, Data Controls — turn off "Improve the model for everyone." In Gemini, turn off Gemini Apps Activity. Claude does not train on your conversations by default for consumer accounts. Business and enterprise tiers of all three carry contractual protection that consumer tiers do not, which is the actual reason to buy one.

Or keep it entirely on your own machine. This is more approachable than it sounds. Download **Ollama** or **LM Studio** — both free — and you can run a capable open model entirely on your own laptop with no internet connection and nothing leaving the building. A Mac with 16GB of memory or a recent Windows machine handles the smaller models comfortably. They are noticeably less capable than the frontier cloud models, and for summarizing, drafting, and reformatting your own documents that gap rarely matters.

The rule that matters regardless. Anyone bound by a confidentiality obligation — attorneys, accountants, physicians, financial advisers, anyone holding client material — should treat a consumer AI tool as a disclosure to a third party, because that is what it is. De-identify before pasting, or use a business tier with a contract behind it.

Two things that are already regulated

These are the exposures I would want you to know about before you leave, because both are live now and neither is obvious.

If you advertise property, algorithmic targeting is a fair housing question. HUD's guidance from May 2024 covers both advertising and tenant screening, and the essential point is that **the advertiser remains responsible**. You cannot delegate compliance to a platform's audience algorithm, and "the tool chose who saw it" is not a defense. Separately, if you enhance listing photographs with AI, California now requires a conspicuous disclosure plus access to the original unaltered image — and other states tend to follow California here.

If you make outbound calls, AI voices are restricted. The FCC ruled in February 2024 that AI-generated voices fall within the TCPA's restrictions on artificial or prerecorded voices, which means using one on an outbound call requires prior express consent. And if you record or transcribe meetings, several states require every participant to consent, not just you.

The fraud protocol, which costs nothing

The FBI's most recent report attributes roughly \$893 million in adjusted losses to AI-enabled fraud, and separately reports \$7.7 billion in losses among victims aged sixty and older. The mechanism has changed in a way most people have not absorbed.

Voice is no longer evidence of identity. It was, for your entire life, until about two years ago. A recording of a few seconds is now enough to clone a voice convincingly.

So: **agree on a phrase with your family and your staff.** Something a stranger could not guess and that does not appear anywhere online. If anyone calls in distress asking for money, a wire transfer, or a gift card, ask for the phrase. If they cannot produce it, hang up and call the person back on the number you already have.

Tell your parents. Tell your children. Tell whoever handles your wire instructions.

Three moves for whatever you run

If you own or lead a company. Pick the three most repetitive things your team does. Automate the one where nobody would be upset if a machine did it. Leave the other two alone for a month while you watch the first. Then ask the question almost nobody asks: **what did the reclaimed hours get spent on?** If the answer is visible — more clients served, a faster close, a project that had been deferred — you have a return. If nobody can say, you have a cost with a good story attached.

If you sit on a board or advise. The failure mode is not the technology. It is the gap between a tool that works in a demonstration and a process that survives a review. Ask management for one number: cycle time on a single named process, measured the same way before and after. Then ask about rework, because the quiet failure of automation is that throughput improves while quality degrades and the two net to zero.

If you sell, raise money, or represent clients. Use these tools to prepare and never to substitute. Write the first draft, then read it out loud before it goes anywhere — the failure mode is not inaccuracy, it is sameness, and sameness is a competitive problem when your letter sits next to forty others written by the same model. The part of your work that is information was commoditized years ago. The part that is telling someone something they do not want to hear, and being trusted while you do it, just became considerably more valuable.

What I would actually do this week

One paid subscription. Two weeks of daily use before buying anything else. Standing instructions set once. The training toggle turned off. A verification phrase agreed with your family tonight. And one repetitive task automated — the smallest one you can find.

That is the whole program. Everything else is refinement.

The one thing worth remembering

Almost everything we experienced as connection was made partly of cost. The forty-second question you walked down the hall to ask a colleague was never valuable as information. It was the doorway that produced the other four minutes. AI answered the question perfectly and removed the doorway.

So the decision in front of all of us is not which tools to adopt. It is which conversations must stay effortful, and protecting those on purpose.

AI can set more places at the table than any of us could set alone. What it cannot do is be a guest.

The full eighteen-page guide — tools, pricing, privacy settings, local installation, and sections for owners, boards, real estate, fundraising, sales and education — is free at DahlingConsulting.com/ai.

If something here touches a decision you are actually facing, write to me and describe it in a paragraph. I will tell you what I think, including when the answer is that you do not need help.

Alicia Dahling, MBA Founder & CEO, Dahling Consulting Alicia@DahlingConsulting.com
DahlingConsulting.com · UnapologeticallyDahling.com · DahlingDigital.com

Sources for every claim in this guide are cited in the full version. Principal references: Cheng et al., "Sycophantic AI decreases prosocial intentions and promotes dependence," Science, March 2026; HUD Office of Fair Housing guidance on digital advertising, May 2024; FCC Declaratory Ruling 24-17, February 2024; FBI IC3 2025 Annual Report, April 2026; California AB 723.