Putting AI to Work

A Working Guide for People Who Have Real Jobs

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

Everything in this guide is something you can do yourself, this week, without hiring anyone. Every price was checked in August 2026 and every one of them will change, so confirm before you buy. Where a claim comes from research or a regulator, the source is named so you can check it rather than take my word for it.

A word on what is deliberately absent. You will not find model version numbers here. They change monthly, and a guide that names them is wrong by October. You will also not find any tool I could not verify from the vendor's own documentation. Several widely recommended products did not survive that check and were removed.

Read the first two sections in order. After that, skip to the section for your role and the pages that interest you.

Part One: Choosing Your Tools

The one thing to understand first

There are two different kinds of AI product, and most people's frustration comes from using the wrong kind for the job.

An **assistant** — ChatGPT, Claude, Copilot, Gemini — generates language. It writes, drafts, summarizes, restructures, and works over documents you give it. Its knowledge came from training, so it can be confidently out of date, and it will occasionally invent things.

An **answer engine** — Perplexity is the clearest example — searches the live web first, then summarizes what it found, with citations you can click. It is not better or worse. It is for a different question.

The practical rule: if the answer needs to be current or sourced, use an answer engine. If the work needs to be written, use an assistant. Most professionals should have one of each.

What each of the major tools is genuinely good at

Tool	Where it earns its keep	Entry price
ChatGPT	The broadest feature set. Best for people who want one tool that does research, drafting, images, scheduled tasks, and simple agents.	Free; Go $8/mo; Plus $20/mo
Claude	Long-document work and writing that does not sound like a machine wrote it. **Its consumer default is not to train on your conversations** — the strongest privacy posture of the major assistants.	Free; Pro $20/mo
Gemini	Already inside Gmail, Docs, Sheets and Drive if you use Google. NotebookLM alone justifies it.	Free; AI Plus $4.99/mo; AI Pro $19.99/mo
Microsoft Copilot	Already inside Word, Excel, Outlook and Teams. Now bundled into Microsoft 365 Personal and Family.	Bundled; Copilot Pro $20/mo
Perplexity	Research where you need to see and click the source. Market scans, competitive checks, fact verification.	Free; Pro $20/mo
Grok	Conversational assistant tied to X. Note its default is to use your data for training.	Tiers via X subscriptions
DeepSeek, Qwen, Kimi	Capable assistants developed outside the US. Free tiers are generous. Consider carefully where the data goes before putting business information into any of them.	Free tiers available
Local models (Llama, Mistral, Gemma, Phi, Qwen)	Work that must never leave your machine. Covered in Part Three.	Free

If you want one recommendation rather than a menu: **pick Claude or ChatGPT as your writing assistant, add Perplexity for research, and use whichever of Gemini or Copilot is already inside the software you own.** That is a complete setup for about forty dollars a month, and two of the three have usable free tiers.

A note on the non-US assistants

DeepSeek, Alibaba's Qwen, and Moonshot's Kimi are genuinely capable and often free. Some of their models are also open-weight, which is why they appear again in the local-AI section — you can run them on your own hardware, which resolves the data question entirely.

Used through their own apps and websites, however, your data goes to servers governed by another country's law. For a recipe or a travel itinerary, that may not concern you. For anything involving clients, employees, contracts, or business strategy, I would not use them, and many organizations prohibit them outright. Check your own policy before you install anything.

Setting up properly, once

Most people use these tools badly because they never spent twenty minutes on setup. These four steps make every subsequent conversation better.

Step one: turn off training, if that matters to you. Instructions for every major service are in Part Two. Do this first, before you paste anything real.

Step two: tell it who you are. Every major assistant has a settings field for persistent instructions. Fill it in once and it applies to everything you ever do.

In ChatGPT: profile icon, then **Settings**, then **Personalization**, then **Custom Instructions**. In Claude, the equivalent lives in your profile settings and in each Project. Write something like this:

> I own a twelve-person consulting firm in Colorado. I value directness and brevity. Do not open your responses with compliments or restate my question back to me. 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.

That last instruction alone will improve your results more than any prompting trick.

Step three: create a Project for each thing you work on repeatedly. A Project is a folder that holds instructions and reference documents so you stop re-explaining context every time. In ChatGPT, click **New project** in the left sidebar. In Claude, click **Projects**, then **New Project**, and add files under **Project knowledge**.

If you are a broker, make one Project per active listing. A fundraiser, one per campaign. A board member, one per organization, with the bylaws and last four board packets loaded in.

Step four: learn one keyboard habit. Most people type one line and accept whatever comes back. Instead, get in the habit of adding two sentences of context and one constraint. Compare "write a follow-up email" against "write a follow-up email to a prospect who raised a pricing objection on Tuesday; under 150 words; no jargon; do not apologize." The second takes nine extra seconds and saves you a rewrite.

The prompt template that actually works

Ignore most of what you have read about prompt engineering. When researchers tested assigning personas to AI — "act as an expert attorney" — across 2,410 factual questions and four model families, it did not improve accuracy at all compared with no persona, and badly chosen personas made results worse [1]. Personas do help with tone and voice. They do not confer expertise.

What does work is structure, examples, and asking for reasoning. Use this shape:

> **Context:** I am a [role] at a [type of organization]. The situation is [two or three sentences of specifics].
>
> **Goal:** [Exactly what you want produced, including length and format.]
>
> **Constraints:** 1. Do not invent facts, statistics, names, or citations. Use only what I have given you. 2. If you need information I have not provided, ask me rather than assuming. 3. [Your own constraint — a rule, a legal requirement, a tone.]
>
> **Before you answer:** think it through step by step, then give me the final version.

That last line is not decoration. Asking a model to work through its reasoning measurably improves accuracy on problems with multiple steps, and it lets you see where it went wrong.

Three prompts that make AI argue with you

Across eleven leading models, researchers found AI affirms the user's actions about forty-nine percent more often than another person would — and it held even when the described conduct involved deception or harm. People who received the agreeable version became less willing to repair conflicts and more convinced they were right [2].

You cannot fix that by hoping. You have to ask for the opposite, deliberately. Keep these three somewhere you can paste them.

> **The critic.** "Here is my plan: [plan]. Act as a harsh, objective critic. Do not flatter me and do not agree with my premise. Identify at least three significant flaws, risks, or logical gaps, and explain concretely why this could fail."
>
> **The opposite case.** "Here is my argument: [argument]. Argue the exact opposite position as persuasively as you can, using evidence and logic. Do not look for middle ground and do not validate my original view."
>
> **The assumption audit.** "Review this plan: [plan]. List every assumption I am making, including the ones I have not stated. For each, explain why it might be wrong and what happens if it is."

Run the critic prompt on any decision that matters before you commit to it. It is the single highest-value habit in this guide.

Verifying before you rely on it

Five checks, none of which take long, and each of which has saved someone from public embarrassment.

Every citation gets clicked. Fabricated sources are the most consequential failure mode of these tools. Attorneys in Mata v. Avianca filed a brief containing cases that did not exist; when they asked ChatGPT whether the cases were real, it confidently confirmed that they were. The court sanctioned them [3]. In Park v. Kim, the Second Circuit referred an attorney to its grievance panel for the same thing [4]. Michael Cohen supplied his own lawyer with fake citations generated by an AI assistant [5]. These were not careless people using a toy. They were professionals who trusted a fluent answer.

Every number gets recomputed. These are language models, not calculators. They will produce a confident total that is wrong. OpenAI's own published explanation of why models hallucinate points at the incentive: evaluation rewards guessing over admitting uncertainty [6].

Anything that sounds exactly like what you hoped gets a second look. That is what sycophancy feels like from the inside.

Negative instructions get checked. If you said "do not mention price," confirm it did not mention price.

Nothing confidential goes in without a decision. If you would not email it to someone outside your organization, do not paste it into a consumer AI tool.

Part Two: Privacy, and How to Opt Out

This is the section most people skip and later wish they had not. It matters for two separate reasons: whether your words are used to train someone's model, and whether you have professional obligations that make that a problem.

What the defaults actually are

Most consumer AI tools use your conversations for training unless you turn it off. The notable exception is Claude, where consumer training is opt-in and off by default.

Service	Default	Where the switch is
ChatGPT	Training on	Profile icon → **Settings** → **Data Controls** → turn off **"Improve the model for everyone"**
Claude	Training **off**	**Settings** → **Privacy** → leave **"Help Improve our AI models"** off
Gemini	Training on	gemini.google.com → **Settings and help** → **Activity** → turn off **"Gemini Apps Activity"**
Copilot (web)	Training on	copilot.com → profile icon → your name → **Privacy** → turn off training on conversation and voice activity
Perplexity	Training on	Profile → **Settings** → **Account** → turn off **"AI Data Retention"**
Grok / X	Training on	**Settings** → **Privacy and safety** → **Data sharing and personalization** → **Grok** → uncheck
LinkedIn	Training on	**Settings & Privacy** → **Data Privacy** → **"Data for Generative AI Improvement"** → off
Meta AI	Training on	Privacy Center objection form; process differs by country

The LinkedIn one surprises almost everyone. It was switched on by default in many jurisdictions with little fanfare, and it covers your profile and posts. It takes about fifteen seconds to turn off.

Four things to understand about opting out

It is forward-looking only. Turning the switch off stops future use. It does not retrieve anything already absorbed. If you have been pasting client material into a consumer tool for two years, that has happened.

Deleting is slower than it looks. OpenAI removes deleted conversations from its systems within thirty days absent a legal hold. Archiving a chat is not deleting it.

Temporary modes are genuinely useful. ChatGPT's **Temporary Chat**, reachable from the model dropdown, is not used for training and is deleted within thirty days. Claude's **Incognito** chat, the ghost icon, is never used for training regardless of your other settings. Use these for anything sensitive that you do not need to keep.

Business tiers are different, and this is the cheapest fix available. Every major vendor excludes business and enterprise accounts from training by default. If you run a company and your people are using AI on client work, moving them onto a business tier costs about twenty-

five dollars per person per month and resolves most of the exposure. That is a smaller number than one afternoon of a lawyer's time.

If you hold a professional license

Attorneys, accountants, physicians, financial advisers and anyone else bound by confidentiality obligations should treat consumer AI tools as a disclosure to a third party, because that is what they are. Entering client-identifying material into a tool that trains on inputs is difficult to defend.

Two practical paths. **De-identify before you paste** — strip names, account numbers and anything uniquely identifying, and ask your question about "a client" rather than about Mr. Harrison. Or **use a business tier with contractual protection.** Zero Data Retention arrangements exist, but they are generally enterprise or API agreements, not something you can select on a consumer signup page.

Moving your email off the major providers

Some people in this room have decided the answer is to stop feeding the large platforms altogether. That is a coherent position, and here is what it actually involves.

Proton Mail, based in Switzerland, is the most common destination. Mail Plus runs about $3.99 a month billed annually, with an Unlimited tier around $9.99 that adds storage, VPN and Drive. It provides zero-access encryption, meaning Proton cannot read your stored mail. To migrate, create the account, then go to **Settings**, then **Import via Easy Switch**, choose Google, authorize, and select mail, contacts and calendars.

Tuta, based in Germany, encrypts more of the mailbox than most competitors, including subject lines and your contact list. Its entry tier runs about three euros a month.

Fastmail, based in Australia, is fast and privacy-respecting but does not offer strict end-to-end encryption. Around five dollars a month.

Two honest caveats. **End-to-end encryption protects the contents of a message, not the metadata** — who wrote to whom, when, and how often remains visible, and metadata is frequently more revealing than content. And **forwarding your Gmail to an encrypted provider accomplishes nothing**, because the mail still lands at Google first. The real work of leaving is updating your address at every bank, brokerage, utility, airline and service you have ever registered with. Budget a weekend and do it in order of importance.

Part Three: Running AI on Your Own Machine

If you do not want your work leaving your building, you can run capable AI models entirely on your own laptop. No account, no internet connection required once installed, no terms of service. This is the section that surprises people most, because it is genuinely free and genuinely private.

The honest tradeoff

A model running on your laptop is meaningfully less capable than the best cloud systems at hard reasoning and very long documents. Your laptop is not a data center. It will also make your fans run and your battery drain.

For a large share of ordinary professional work — drafting, rewriting, summarizing, brainstorming, extracting information from a document, cleaning up dictation — a local model is entirely sufficient. And for material you cannot risk sending anywhere, sufficient and private beats excellent and exposed.

What your hardware can handle

Model size is measured in parameters, quoted in billions. More parameters means more capable and more memory required. What governs your options is RAM.

Your machine	What to run	Realistic experience
8 GB RAM	The smallest models only	Works, noticeably slow, fine for drafting and summarizing
16 GB RAM	Mid-size models — the sweet spot	Genuinely useful for daily work
32 GB or more	Larger models	Approaches cloud quality for many tasks

One important point in Apple's favor: Macs with M-series chips use unified memory, shared between processor and graphics. **A 16 GB MacBook will generally outperform a 16 GB Windows laptop** for this purpose, unless that Windows machine has a strong discrete NVIDIA graphics card.

Which model to download

I am deliberately not naming versions, because new ones ship constantly and any version I print here will be stale within weeks. Instead, learn the families and pick the newest one that fits your memory.

Look for **Llama** from Meta, **Mistral** from the French lab of the same name, **Gemma** from Google, **Phi** from Microsoft — which is specifically designed to be small and efficient, making it the right first choice on an 8 GB machine — and **Qwen** or **DeepSeek**, both strong and both available as open weights.

When you search inside the tool, you will see several versions of the same model with labels like Q4 or Q5. Those are compression levels. **Choose something in the Q4 or Q5 range**: it is the standard balance of speed against quality, and it is what nearly everyone should use.

Installing LM Studio, step by step

This is the friendliest option and looks much like ChatGPT once running.

1. Go to **lmstudio.ai** and click the download button for your operating system. The application is roughly 400 MB.
2. Install it the way you would install any other program.

3. Open LM Studio. Use the **search** field at the top to look for one of the model families above.
4. From the results, pick a version whose file size is comfortably smaller than your available RAM, and whose name includes **Q4** or **Q5**. Click **Download**. Expect several gigabytes and several minutes.
5. When it finishes, click the **Chat** icon in the left sidebar.
6. At the top of the chat window, open the model selector and choose the model you just downloaded. Wait a few seconds while it loads into memory.
7. Type. You are now running AI with no internet connection and no account.

Ollama, at **ollama.com**, is the other well-established option. It runs quietly in the background rather than presenting a window of its own, which makes it the better choice if you want other applications to use your local model. Both are free for local use.

Chatting with your own documents, privately

The application most people want is asking questions of their own files — contracts, board packets, medical records, tax documents — without those files going anywhere.

AnythingLLM is the approachable tool for this. It creates workspaces, you drop documents in, and you ask questions answered from those documents. Install it, follow its setup wizard, and when it asks for a model provider, point it at LM Studio or Ollama if you have one running. Then create a workspace, upload your files, and use the **Move to Workspace** and **Save and Embed** commands to index them. After that, ask away.

One caution worth stating plainly: running a model locally means your text is not sent to a cloud service, which is a real and significant protection. It does not make your laptop secure in every other respect. Ordinary discipline still applies — full-disk encryption, a strong password, current updates, and care about what you download.

On your phone

Apple Intelligence processes many requests directly on the device. When a request needs more capability, Apple routes it to **Private Cloud Compute**, which the company states does not retain your data or use it for training [7]. **Google AI Edge** brings small on-device models to Android, letting some features work without sending data anywhere [8].

Neither is as capable as a full cloud assistant. Both are meaningfully more private, and both are already on the phone in your pocket.

Part Four: Building Something That Runs Without You

Everything so far has been a conversation you start. This is about work that happens whether or not you are at your desk. You do not need to write code for any of it.

Start with the boring wins

The highest-return automation is almost never impressive. It is the form that fills in your CRM by itself, the invoice reminder that goes out on day thirty-one, the meeting notes that arrive as action items instead of a transcript.

Zapier connects several thousand applications and has a free tier covering a hundred tasks a month, with paid plans from about twenty dollars. **Make** does the same thing with a more visual interface, starting around eleven dollars. Either will let you build, in under an hour, a chain like: when someone submits the contact form, create a CRM record, send me a summary, and add a follow-up task for Thursday.

For meetings, connect a transcription tool to your calendar and stop taking notes. Otter has a free tier with paid plans from about ten dollars per user. Microsoft Teams and Google Meet both do this natively if you already pay for them. Then use one prompt on the transcript:

> "From this transcript, extract: every decision made, every action item with the person responsible and the deadline, any question left unanswered, and a three-sentence summary. Note anything where the commitment was ambiguous."

That last clause is the valuable one. Ambiguous commitments are where projects die.

Building an actual agent, in an afternoon

MindStudio at mindstudio.ai is the most approachable platform for building a purpose-built AI tool without programming. The free tier includes one agent and a thousand runs a month; the individual plan runs about twenty dollars for unlimited agents. It connects to a wide range of underlying models without your needing to manage keys, and you can upload your own documents so the agent answers from your material rather than general knowledge.

A realistic first project, start to finish:

1. Sign in at mindstudio.ai and choose to create a new agent, starting from scratch.
2. Name it for the job you want done — "Listing Description Drafter," "Grant Narrative Reviewer," "Board Packet Summarizer."
3. In the instructions or system prompt field, write the standing brief. Be specific about what it must never do. For example: You draft MLS listing descriptions from the features I provide. You never invent features. You never use language implying a preference about who should live in a home, including phrases like "perfect for a family." You return exactly 150 words.
4. Add an input field so you can paste in the raw material each time.
5. Upload any reference documents you want it to follow — your style guide, your compliance checklist, three examples of listings you were happy with.
6. Test it with real material. Then adjust the instructions where it got things wrong, which is where the actual work is.
7. Publish it, and use the link. You now have a tool that behaves consistently instead of a chat you have to re-explain every time.

The other platforms worth knowing: **Microsoft Copilot Studio** if your organization already lives inside Microsoft 365, and **n8n**, which is powerful and can be self-hosted — though I would not recommend a non-technical person attempt to self-host anything; use the cloud version.

Agents that browse and act, and the reason to be careful

The genuine change of the past year is that these systems moved from answering to acting. ChatGPT's agent capability, available on paid consumer plans, operates its own virtual computer: it navigates websites, fills forms, runs code, and produces files. It asks permission before consequential actions.

Before you hand anything real to an agent, understand **prompt injection**. Hidden instructions can be embedded in a web page or a document — invisible to you, perfectly legible to the agent that reads it. A compromised page can instruct an agent to exfiltrate what it has access to or take actions you never authorized [9]. This is not theoretical and it is not solved.

Four rules I follow and recommend without exception. Never give an agent unsupervised access to a payment method or a financial account. Never let it send messages on your behalf without your reading them. Do not let it browse unknown sites while it holds access to anything sensitive. And keep a human in the loop for anything consequential — which, restated, is simply the earlier rule: **supervision proportional to the cost of being wrong.**

Presentations, images, and video

Gamma at gamma.app turns an outline into a presentation or one-pager. There is a free tier with one-time credits; paid plans run about nine to eighteen dollars a month. It produces a usable first draft in a minute and requires real editing afterward, and it will occasionally state something that is not true. Treat it as a fast layout engine, not an author.

Canva has its AI features bundled into Pro, around a hundred and twenty dollars a year. **PowerPoint** will build a deck from a Word document if you have Copilot. **Google Slides** will generate visuals in place.

For images and video, the question that matters is not quality but whether you may legally use the output.

Adobe Firefly is the commercially safest choice, trained on Adobe Stock, openly licensed and public domain material, with non-beta output cleared for commercial use. Standard is about $9.99 a month. **Midjourney** produces the most striking images, from about ten dollars a month — but note carefully that **if your business grosses over one million dollars a year you must be on its Pro or Mega tier to use output commercially** [10]. **Runway** for video, from about twelve dollars, with commercial rights retained by you. **ElevenLabs** for voice, with a commercial license in all paid plans and a firm prohibition on cloning a voice without permission. **Higgsfield** at higgsfield.ai is a more advanced creative suite offering camera, lens and lighting control, from about nineteen dollars a month, aimed at people producing serious visual work.

Four situations where you should not use AI imagery at all. **Real estate photography that misrepresents a property** — covered specifically in the next section. **Journalism. Any testimonial or endorsement**, because the FTC's rule banning fake reviews explicitly covers AI-generated ones [11]. And **anything implying a real person said or did something they did not.**

If you publish AI-generated media, know that Meta, TikTok and YouTube all require disclosure of realistic synthetic content, and that the **C2PA Content Credentials** standard now embeds provenance information into files as a kind of nutrition label [12].

Part Five: Your Role

If you own or run a business

Start from the most useful finding in the field. MIT's study of enterprise deployments found that **ninety-five percent produced no measurable impact on the profit and loss statement** — and that **buying a tool built for a specific workflow succeeded about two-thirds of the time, while building internally succeeded roughly a third as often** [13].

So: **buy, do not build.** I say that as someone who builds these systems professionally. Unless you are in the software business, a custom build is where good intentions and real money go to die.

Deal with shadow AI, because it is already happening. A study reported this spring found forty-nine percent of workers using AI in ways their employer had not approved, most of them on free tiers with no data governance, and eighty-six percent of IT leaders reporting at least one negative incident tied to unapproved use [14]. Your people are not being insubordinate. They are being resourceful, and the official option was worse than the unofficial one.

Find out what they are using with a message that carries no threat in it:

> "We are taking stock of our software. Please reply with any AI tools you currently use for work — ChatGPT, Grammarly, Canva, anything. This is not a compliance exercise; I want to know what is genuinely helping so we can pay for the good ones properly and make sure nobody is putting client information somewhere it should not go."

Then write the policy. Five lines is enough:

1. Treat AI like sending information outside the company. If you would not email it to a stranger, do not paste it.
2. No client names, account numbers, employee matters, or confidential financials in public tools.
3. AI produces drafts. A person owns the final version and is accountable for it.
4. Use the tools we pay for. Ask before adding a new one.
5. When in doubt, ask. Nobody is in trouble for asking.

Measure the right thing. Most AI measurement I see is activity dressed as outcome: seats licensed, prompts entered, hours "saved." Hours saved is meaningless unless you can name what filled them. Ask instead: what did we stop doing, where did the freed capacity actually go, and what got worse? Something always gets worse, and if nobody can tell you what, it means nobody looked.

And watch where the friction went, not only where the cost went. If you saved two hundred thousand dollars and your twenty-somethings stopped talking to each other, you did not save two hundred thousand dollars. You moved a cost onto a line that does not appear on any report you receive, and it returns in about eighteen months as turnover you will attribute to something else.

Five starting projects, in the order I would do them: automate one form-to-system handoff; put transcription on every recurring meeting and extract action items; draft first versions of

routine customer communications; handle your five most repetitive customer questions with a chat tool on your site; and write down the decisions in your organization that will never be automated.

If you are in real estate

You are among the highest adopters in any profession — the National Association of Realtors found sixty-eight percent of agents using AI, forty-six percent for content such as listing descriptions [15] — and you are also among the most regulated. The compliance material below is the part worth more than the prompts.

Fair housing is the real risk, and AI walks into it naturally. HUD has issued explicit guidance that algorithmic ad targeting and delivery can unlawfully deny people information about housing based on protected characteristics, and that the advertiser remains responsible [16]. AI-written copy drifts into this without any intent on your part: "perfect for a family," "bachelor pad," "exclusive neighborhood," "safe area," "walking distance" — each implies a preference or makes an assumption about who belongs. Instruct the tool explicitly and check every draft.

A prompt with the guardrail built in:

> "Write a 150-word MLS listing description for a three-bedroom, two-bath home. The features are: [list them]. Strict requirements: comply with the Fair Housing Act. Do not use any language suggesting a preference or limitation based on race, color, religion, sex, disability, familial status, or national origin. Do not use 'perfect for a family,' 'bachelor,' 'exclusive,' 'safe,' or 'walking distance.' Describe only the physical property and its objective location. Do not invent any feature I have not listed."

Photography rules have tightened significantly. California's AB 723, effective January 1, 2026, requires clear written disclosure for any digitally altered or virtually staged listing photo, with the unaltered original made available. MLS practice generally permits virtually adding or removing personal property — furniture, decor — but prohibits altering or concealing permanent features such as walls, flooring, power lines, or neighboring structures. NAR's Code of Ethics Article 12 requires a "true picture" in advertising, and that applies to AI-generated copy and imagery alike [17].

Before any photo goes up, confirm four things: whether it was altered at all; that no permanent feature was changed or hidden; that the image is labeled if it was staged; and that you have the original available. Then put the disclosure in the public remarks.

Where AI genuinely helps you: comparable-sales research and CMA preparation, though the valuation judgment remains yours; drafting and rewriting listing copy for different channels; follow-up sequences so leads do not go cold; transaction coordination checklists; and turning one property into a week of social content. What it should not do is determine value, give legal advice, or touch client financial details in a consumer tool.

NAR now advises every brokerage to adopt a formal AI use policy [18]. If yours has not, the five-line policy above is a serviceable starting point with fair housing added as a sixth line.

If you raise money for a nonprofit

The sector-wide finding should shape how you spend your effort. Ninety-two percent of nonprofits now use AI, and **seven percent report a major improvement in their ability**

to achieve their mission. Eighty-one percent use it ad hoc and forty-seven percent have no governance policy at all [19]. Researchers call it an efficiency plateau: everyone is drafting faster and almost nobody has changed what the organization can do.

So the question is not whether you use AI. It is **what you stopped doing with the time**, and whether that capacity went into donor relationships or simply evaporated.

On grant applications, one finding should govern your behavior. Candid found that fifty-seven percent of grantmakers do not know whether they have received AI-generated applications because applicants rarely disclose, only ten percent explicitly accept AI-generated content, and sixty-seven percent have no policy at all [20]. That is not permission. That is an unsettled question you should not resolve unilaterally in your own favor.

The defensible line: **use AI to polish language, never to author your program logic or theory of change.** Funders are buying your judgment about how change happens in your community. And verify every statistic and citation by hand, because a fabricated source in a proposal is disqualifying and memorable.

For appeals, the tool is genuinely useful — with a rule. Never put real donor or beneficiary information into a consumer tool. Write with composites:

> "Draft a 300-word email appeal for our spring campaign. We provide after-school tutoring for middle school students. The goal is $5,000 for tablets. Use a composite, clearly hypothetical student example — do not use or invent a real name. Tone: urgent but hopeful, not maudlin. End with a specific ask for $50. No exclamation points."

Sector guidance is converging on the same points: the Association of Fundraising Professionals advises getting informed consent before using supporter data with AI tools and publishing your AI policy publicly, and the BBB Wise Giving Alliance advises setting boundaries before deployment and maintaining meaningful human oversight [21] [22]. Publishing a short AI policy on your website is becoming a marker of governance maturity, and it takes an afternoon.

If you are in sales or business development

The legal constraints are where I would spend your attention, because they carry real penalties and most people do not know them.

AI voices on outbound calls. The FCC ruled in February 2024 that AI-generated voices in calls are "artificial" under the Telephone Consumer Protection Act, which makes using them for telemarketing without prior express written consent illegal [23]. This is not a gray area.

AI meeting recorders and consent. Transcription bots are governed by state wiretapping law. In all-party consent states — California, Connecticut, Florida, Illinois, Maryland, Massachusetts, Michigan, Montana, Nevada, New Hampshire, Pennsylvania and Washington among them — every participant must agree. If you sell across state lines, assume the strictest rule applies, announce the recording, and get an audible yes.

Email deliverability. Google, Yahoo and Microsoft require bulk senders to keep spam complaints below 0.3 percent and bounces under 2 percent, with SPF, DKIM and DMARC authentication in place. Volume is what gets you filtered, and no amount of clever copy survives a poor sender reputation.

Where AI genuinely earns its place in a sales week: **call preparation**, which Gartner expects to be the origin of ninety-five percent of seller research workflows by 2027, up from under twenty percent in 2024 [24]; **CRM hygiene**, converting a transcript into structured notes; **follow-up drafting** that references what was actually said; and **objection rehearsal**, which is the most underused application of all.

> "I am presenting a proposal to a mid-market manufacturer. We are roughly fifteen percent above our main competitor, and we differentiate on [X] and [Y]. Act as a skeptical CFO who does not want to switch vendors. Give me your three strongest objections, then a concise value-based response to each. Do not be gentle with me."

The failure mode to avoid is using AI to increase volume. Buyers can tell, and the market has already adjusted: personalization is now the only thing that works, and mass AI-generated outreach damages both your reply rate and your sender reputation. Use it to know more about fewer people.

If you are a retired executive or serve on boards

Two very different things belong in this section, and the second one is more urgent than the first.

The genuinely useful applications. For board work, the single best tool available is **NotebookLM**, free with a Google account at notebooklm.google.com. Upload the board packet, the bylaws, the last four sets of minutes, the audit — then ask questions answered strictly from those documents, which sharply reduces invention. Its **Audio Overview** feature turns your own documents into a two-host discussion you can listen to on a drive, which is the most useful thing I have found for a two-hundred-page packet the night before a meeting.

For minutes and summaries, this prompt does most of the work:

> "Acting as a corporate secretary, review this transcript. Provide a concise executive summary, the decisions actually made as distinct from matters merely discussed, a list of action items with owners and deadlines, and any question raised but left unresolved. Flag anything where the record is ambiguous about what was agreed."

Note firmly: **AI-drafted minutes are a draft.** They are not the record until a human has verified them, and the official record has legal weight.

Beyond the boardroom, these tools are excellent at travel itineraries with real constraints ("relaxed pacing, no more than two activities a day, narrated tours to limit walking, centrally located and accessible hotels, quieter restaurants for easier conversation"), at helping you prepare questions before a medical appointment, at learning a subject from scratch, and at researching organizations before you give to them. Any assistant with voice input will also conduct a guided life-story interview if you ask it to — you talk, it organizes, and you have the beginning of a memoir your family will actually read.

Now the part I would insist on if I could. AI-enabled fraud is targeting this demographic specifically and effectively. The FTC reports fraud losses among people over sixty rising from about six hundred million dollars in 2020 to two point four billion in 2024, with investment scams the costliest category and a hundred fifty-nine million lost to tech-support scams in a single year [25]. The FBI's Internet Crime Complaint Center reports more than two hundred

one thousand victims aged sixty and older reporting over seven point seven billion dollars in losses, roughly thirty-seven percent higher than the year before [26].

The mechanism that matters most is voice cloning. A few seconds of audio — a voicemail greeting, a video posted at a graduation — is enough to produce a convincing imitation of your grandchild's voice, calling in distress and asking for money urgently.

The protocol, and it costs nothing:

Agree with your family on a **phrase** that no machine could infer from anything any of you has ever posted. Not a pet's name, not a hometown, not a birthday. If someone calls in distress asking for money, you ask for the phrase before anything else.

Then, regardless of the answer: **hang up and call back** on the number you already have saved. If you cannot reach them, call another family member. And treat any request for a wire transfer, cryptocurrency, or gift cards as conclusive evidence of fraud, because legitimate emergencies do not require untraceable payment [27]. Report anything suspicious at ReportFraud.ftc.gov.

The FBI's cyber division has said plainly that these can appear legitimate to even the most trained individuals. **The voice is no longer evidence.** It was for our entire lives, and it is not anymore.

If you teach, lead a school, or have children at home

The upside is real and measured. A randomized trial in a Harvard undergraduate physics course found AI tutoring outperformed in-class active learning [28]. A World Bank evaluation in Nigeria found six weeks of after-school AI tutoring raised test scores meaningfully [29]. Read the second study's caveat carefully, though: it was **teacher-led**, and the authors attribute the result to teacher facilitation. The AI did not replace the teacher. It worked because there was one.

Where it saves real time: rewriting a text at a different reading level, drafting rubrics, generating differentiated versions of an activity, drafting parent communications, and preparing first drafts of documentation. Two that work immediately:

> "Rewrite this fifth-grade science passage at a third-grade reading level, preserving the core concepts of evaporation, condensation and precipitation, using simpler vocabulary and shorter sentences. Then write three comprehension questions on the simplified version."
>
> "Create a four-point rubric for a seventh-grade persuasive essay assessing thesis clarity, use of evidence, organization and transitions, and grammar. Format as a table."

The hard constraints. Student personally identifiable information must not go into consumer AI tools that lack a data privacy agreement with your school — that is FERPA, and COPPA and state student-privacy laws add to it. Use initials or describe the situation generically. If you are a verified US K-12 educator, note that **ChatGPT for Teachers** is free through June 2027 and does not train on your data, which resolves much of this [30].

On AI detection, be careful. Stanford researchers found detection tools are biased against non-native English writers, with detectors unanimously flagging nineteen percent of essays by non-native writers as AI-generated [31]. A false accusation of cheating is a serious harm to a student. Detectors should be one input among several — alongside document version history and an actual conversation with the student — never the sole basis for discipline.

For parents and grandparents, the most useful thing is not surveillance. It is asking, without judgment in your voice, **"what's it good for?"** rather than "how much time are you spending on that." The second question ends the conversation; the first one starts it. Three questions worth asking a child:

> "When you ask it something, how do you think it comes up with the answer? Does it actually know things, or is it predicting what usually comes next?"
>
> "It told us this. How would we check whether it's true?"
>
> "If it writes the whole essay, what did you learn? Where's the line for you?"

Worry much less about whether they use it and much more about what it is replacing. If it is replacing boredom, that is fine. If it is replacing you, that is information.

The Three Questions

Everything in this guide comes down to three questions. Pick one. You do not need all three.

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

You will know the answer before you finish asking, which is exactly what makes the question useful. The distinction belongs to Simone Heng, who studies human connection: is this a **way station** or a **destination**? A way station is somewhere you pass through on the way to a person. A destination is where you stop. The person who rehearses a hard conversation with AI and then goes and has it does better. The person who has it with the AI instead does worse. Same tool, same afternoon, opposite outcome — and the technology did not make that choice.

Did the effort matter here?

Usually not. Let the machine draft the agenda, summarize the packet, clean up the memo. But for a condolence note, an apology, or telling someone you are proud of them, the effort is the message. AI can give you better words in four seconds. It cannot give away your twenty minutes, and the twenty minutes was the entire gift.

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. Go back and run the critic prompt.

If you want to go further

BlueDot Impact at bluedot.org is a London non-profit that grew out of a Cambridge student reading group and has since taken roughly six thousand people through courses on AI safety and strategy, a number of whom now work at the frontier labs and at the UK's AI Security Institute. The courses are free on a pay-what-you-want basis; the AGI Strategy course runs about thirty hours. They are a safety organization with a point of view, which is worth knowing going in.

I applied a few weeks ago, and the reason is the syllabus. I expected a course on AI risk to open with threats. Unit One is called Imagining a Better Future, and before you analyze a single risk you are asked to write down the future you actually want.

You cannot build a defense for a future you have never described. And a way station only makes sense if you know where you are going.

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

References

[1] PromptHub, "Role-Prompting: Does Adding Personas to Your Prompts Really Make a Difference?" — https://www.prompthub.us/blog/role-prompting-does-adding-personas-to-your-prompts-really-make-a-difference [2] Cheng et al., "Sycophantic AI increases attitude extremity and overconfidence," Science, March 2026 [3] Mata v. Avianca, Inc., S.D.N.Y., 2023 (sanctions for fabricated AI-generated citations) [4] Park v. Kim, 2d Cir., 2024 (referral to grievance panel for citing a non-existent case) [5] Reuters, "Michael Cohen won't face sanctions after generating fake cases with AI," March 20, 2024 — https://www.reuters.com/legal/michael-cohen-wont-face-sanctions-after-generating-fake-cases-with-ai-2024-03-20/ [6] OpenAI, "Why language models hallucinate," 2025 — https://openai.com/index/why-language-models-hallucinate/ [7] Apple, "Private Cloud Compute" — https://security.apple.com/blog/private-cloud-compute/ [8] Google AI Edge — https://developers.google.com/edge [9] Palo Alto Networks Unit 42, "AI Agent Prompt Injection" — https://unit42.paloaltonetworks.com/ai-agent-prompt-injection/ [10] Midjourney, "Comparing Midjourney Plans" — https://docs.midjourney.com/hc/en-us/articles/27870484040333-Comparing-Midjourney-Plans [11] Federal Trade Commission, "FTC Announces Final Rule Banning Fake Reviews and Testimonials," August 2024 — https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials [12] C2PA Content Credentials specification — https://spec.c2pa.org/ [13] MIT Project NANDA, "The GenAI Divide: State of AI in Business 2025" — reported at https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/ [14] Forbes, "Your Employees Are Using AI Without Telling You. Now What?" April 30, 2026 — https://www.forbes.com/sites/terdawn-deboe/2026/04/30/your-employees-are-using-ai-without-telling-you-now-what/ [15] National Association of Realtors, 2025 Technology Survey — https://www.nar.realtor/press-releases/realtors-embrace-ai-digital-tools-to-enhance-client-service-nar-survey-finds [16] U.S. Department of Housing and Urban Development, "HUD Issues Fair Housing Act Guidance on Applications of Artificial Intelligence," May 2, 2024 — https://archives.hud.gov/news/2024/pr24-098.cfm [17] National Association of Realtors, "Are You 'Catfishing' Buyers With Picture-Perfect Real Estate Photos?" February 2026 — https://www.nar.realtor/news/real-estate-news/sales-marketing/are-you-catfishing-buyers-with-picture-perfect-real-estate-photos [18] National Association of Realtors, "Why Every Brokerage Needs an AI Use Policy," March 2026 — https://www.nar.realtor/news/broker-news/why-every-brokerage-needs-an-ai-use-policy [19] Virtuous, 2026 Nonprofit AI Adoption Report, February 2026 — https://virtuous.org/blog/2026-nonprofit-ai-adoption-report/ [20] Candid, "Where do foundations stand on AI-generated grant proposals?" — https://candid.org/blogs/funders-insights-on-ai-generated-grant-application-proposals/ [21] Association of Fundraising Professionals, "How to Leverage AI and Stay True to Your Nonprofit Values" — https://afpglobal.org/how-leverage-ai-and-stay-true-your-nonprofit-values [22] BBB Wise Giving Alliance, "Responsible AI. Built on Trust." — https://give.org/ai [23] Federal Communications Commission, "FCC Makes AI-Generated Voices in Robocalls Illegal," February 8, 2024 — https://www.fcc.gov/document/fcc-makes-ai-generated-voices-robocalls-illegal [24] Gartner, "Use AI in sales to enhance automation, personalization and customer

satisfaction" — https://www.gartner.com/en/sales/topics/sales-ai [25] Federal Trade Commission, "FTC Issues Annual Report to Congress on Agency's Actions to Protect Older Adults," December 1, 2025 — https://www.ftc.gov/news-events/news/press-releases/2025/12/ftc-issues-annual-report-congress-agencys-actions-protect-older-adults [26] FBI Internet Crime Complaint Center, 2025 Annual Report — https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf [27] Federal Trade Commission, "Scammers use AI to enhance their family emergency schemes" — https://consumer.ftc.gov/consumer-alerts/2023/03/scammers-use-ai-enhance-their-family-emergency-schemes [28] Kestin et al., "AI tutoring outperforms in-class active learning," Scientific Reports, 2025 — https://www.nature.com/articles/s41598-025-97652-6 [29] De Simone et al., "From Chalkboards to Chatbots," World Bank, 2025 — https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099548105192529324 [30] OpenAI, "ChatGPT for Teachers" — https://openai.com/index/chatgpt-for-teachers/ [31] Stanford HAI, "AI Detectors Biased Against Non-Native English Writers," 2023 — https://hai.stanford.edu/news/ai-detectors-biased-against-non-native-english-writers

---

Alicia Dahling, MBA Founder & CEO, Dahling Consulting Two decades in corporate finance at HP, Hewlett Packard Enterprise and DXC Technology

DahlingConsulting.com · UnapologeticallyDahling.com · DahlingDigital.com
Alicia@DahlingConsulting.com

Where to put AI — and where to deliberately keep it out.