

THE  
**AI Authorship  
Field Guide**

---

*How to use the new tools without losing yourself*

---

BY

Rebecca Holder

*Real Humankind*

# *Whose story is it anyway?*

— Rebecca Holder

# Contents

<i>One</i>	The Shift You've Been Missing
<i>Two</i>	The Map
<i>Three</i>	How To Actually Use This
<i>Four</i>	What Just Became Possible
<i>Five</i>	What AI Can't Do
<i>Six</i>	The 30-Day Starter Plan

---

<i>A</i>	The Toolbox
<i>B</i>	Prompts To Steal

A NOTE ON THE AUTHOR

---

Rebecca Holder is a writer, founder, and former senior corporate operator. Twenty-five years across Accenture, the BBC, Barclays, and BP. She helped launch BBC iPlayer and Barclays Faster Payments, and her first job was on Absolutely Fabulous.

She is the founder of Real Humankind and writes the Substack *Whose Story Is It Anyway?* on storytelling, identity, and AI.

*A fuller biography sits at the back of this guide, along with all the ways to stay in touch.*

SECTION ONE

## The Shift You've Been Missing

---

*Most of my friends don't get how I talk to my computer all day.*

I'll mention something I worked through with AI — a piece of writing, a difficult email, a pattern I spotted in my own behaviour — and they'll tilt their heads slightly. "What, like ChatGPT?" Yes. "Doing your work for you?" No. Thinking with me.

There's a gap between how AI is being marketed and what it's actually useful for, and that gap is where most people are getting stuck. The marketing is overwhelmingly about productivity. Get more done. Save time. 10x your output. The whole register is efficiency-coded. Faster. Leaner. More.

If you're already, almost certainly, doing too much, the productivity frame is not particularly interesting. You don't need to do more. You probably need to do less, more deliberately, with more of yourself in it.

That's the shift I want this guide to land.

**AI is not, primarily, a productivity tool.**

***It is an authorship tool.***

Productivity is what you do for someone else's agenda — your boss, your clients, your inbox. Authorship is what you do for your own. The difference matters. Most AI writing online assumes you want to produce more for the same agenda you've been serving for years. This guide assumes you want something else: to think more clearly about what your agenda actually is, and to use the tools to take that thinking further than you could on your own.

Here's what AI is genuinely good at, when you use it this way:

- **Spotting patterns you can't see in yourself.** Read it a paragraph of your own writing and ask what you seem to be avoiding. Tell it about a difficult colleague and ask what your default conflict-style is. The model has no investment in protecting your self-image. It will see things you can't.
- **Drafting things you can then edit toward what you actually mean.** Most people use AI for the final draft. Smarter is to use it for the throwaway draft that helps you find your real voice. Read what it produced, notice what's wrong about it, and the wrongness will teach you what you actually want to say.
- **Holding messy thinking long enough for you to see its shape.** You don't have to arrive with polished sentences. Talk to it — literally, by voice — through a tangle of thoughts, and ask it to reflect back the structure underneath. It will be wrong about some of it. The bits it gets right are often the bits you didn't know you were ready to admit.
- **Building things that used to require a team.** A landing page. A first version of an app. A prototype of a business idea. A bot that answers questions about your work. None of this requires you to learn to code anymore. You can describe what you want and the tools will build a version of it. Some of what used to take a startup six months and £100,000 you can now sketch in a weekend on your sofa.
- **Doing the boring work that frees up your real attention.** Restructuring a long document. Summarising twelve articles into one map. Generating fifty title options for a piece of writing so you can pick from them rather than starve in front of a blank page. The boring tasks are real and AI handles them well. But this is the smallest part of what it can do, even though it's the part most articles obsess over.

Notice what is not on this list. Replacing your judgment. Understanding context only you have. Knowing your relationships, your history, your taste, your actual life. That's the territory the technology can't enter — and that's where your value lives.

The question this guide is built around is not how do I use AI to keep up. The question is how do I use AI to take back authorship of work, life, and identity in a moment when most people are using it to do the same things, faster?

If that question interests you, the rest of this is for you.

## SECTION TWO

## The Map

---

Most people, when they hear "AI", think of one tool. Usually ChatGPT. That's like thinking the internet is one website.

The actual landscape has several distinct categories, each doing something quite different. You don't need to use all of them. You probably only need two or three. But knowing what's out there means you can pick the right tool for the right job rather than forcing one to do everything.

### The thinking partners

These are the conversational tools — ChatGPT, Claude, Google's Gemini, Meta's Llama. You talk to them, in text or voice, and they respond. They're general-purpose. They draft, summarise, brainstorm, edit, analyse, plan. Of these, my personal preference is Claude for nuanced thinking and writing, and ChatGPT for everyday tasks — but they're broadly interchangeable for most uses, and the differences between them are less important than the difference between using one and using none.

### The agents

Newer category. These don't just talk to you — they take actions. Claude Code can write and edit code on your computer. Cowork can manage files and tasks across your desktop. Browser-based agents can navigate the web, fill in forms, complete tasks across sites. Think of these as the difference between a colleague who advises you and one who actually goes off and does the work.

### The image-makers

A separate category, though many of the chat tools now include image generation inside them. Midjourney is the most established for high-quality imagery. Nano Banana, DALL-E, Imagen, Stable Diffusion all sit here. You describe what you want to see and the

model makes it. Useful for everything from social media imagery to mock-ups of products that don't exist yet.

## The builders

This is the most under-appreciated category and the one that quietly changes everything. Replit, Lovable, Bolt, v0, Cursor — tools that let you describe a piece of software and have it built. You can make a landing page, a working app, an internal tool, an MVP of a business idea, without writing a line of code yourself. You describe in plain English what you want, the tool builds it, you test it, you refine. Anyone with an idea and a weekend can now build the thing. That hasn't been true at any other point in human history.

## The connectors

Tools like Zapier, Make, n8n — they let different services talk to each other automatically. AI is increasingly built into them. You can set up a workflow where a new email triggers an AI summary which gets posted to your project tool which alerts your phone. The plumbing of automation, increasingly handled by you alone rather than by an IT department.

That's the basic shape. You don't need to learn all of it. For most people, starting with one good thinking partner (Claude or ChatGPT) and one image tool when needed is plenty. The agents and builders are where it gets interesting if you want to make things, not just write things.

The point of this guide is not to teach you all the tools. It's to show you that you can do a remarkable amount with just a thinking partner and an hour a day. Once you understand what's possible, you can decide where to go deeper.

## SECTION THREE

## How To Actually Use This

---

The fastest way to understand what AI can do is to stop reading about it and use it for something that matters to you. So this section is built around two practical demonstrations — one for thinking, one for writing — that you can do today, in fifteen minutes each, with whichever thinking partner you've chosen.

### Demonstration One: The pattern-spotting conversation

Most of us are running silently on patterns we cannot see. Patterns of language, of self-positioning, of small habitual concessions we make without realising. They show up in our writing, our messages, our decisions. Other people sometimes spot them — a good friend, a sharp therapist — but most of the time they go unobserved, including by us.

AI can hold up the mirror.

Here's how it works. Open Claude or ChatGPT. Switch to voice mode if you have it — it's a more natural way to do this. Then say something like:

*"I want to spot some patterns in my own thinking that I can't see for myself. I'm going to read you some recent messages I've written — emails, texts, anything. Don't try to fix them. Just notice what's running through them. Tell me what you see about how I write to people, what I tend to do, what I might be avoiding. Be specific and don't soften it for my benefit."*

Then read it three or four pieces of recent correspondence. A long email to a colleague. A text exchange with a difficult friend. A message to a family member. The more honest the sample, the more useful the analysis.

The model will reflect back things like: You apologise twice in the opening of every message. You frame your requests as questions rather than statements. You consistently describe your own work as

"just" or "only". When pushed back on, you immediately offer to do extra work to compensate.

Some of what it says will be wrong or off. Ignore that. The bits that land are the bits that matter. The hit-rate on this exercise is genuinely uncanny because the model has no investment in your self-image. It is not protecting you. It is just describing the patterns it sees.

Once you have the patterns named, you can ask follow-up questions:

*"Where do you think these patterns might be coming from?"*

*"What might my professional life cost me if these patterns continue?"*

*"What would it look like if I sent these same messages without those patterns? Show me."*

The exercise takes about twenty minutes. Most people I've taken through it report that something quietly shifts. Not a transformation. A noticing. Once you can see the pattern, you can choose whether to keep running it.

## **Demonstration Two: The throwaway-draft technique**

The myth about AI writing is that you should use it to produce your final copy. Get it to write your blog post. Get it to write your LinkedIn update. Get it to write your book.

That approach mostly produces bland writing that sounds vaguely like everyone else's AI output. The reason is that AI is trained on the average of everything ever written, and the average is not interesting. If you ask it for a finished article, it gives you the average article. Useful as a starting point, deadly as a destination.

The smarter use is the opposite: get AI to write the first draft specifically so you can read it and notice what's wrong.

Here's the process. You're writing something — an essay, a pitch, a difficult email, a chapter. Instead of staring at the blank page, you describe what you're trying to say, badly, in voice. Three or four minutes of half-formed sentences. "I want to write something about X. The argument is roughly that Y. I'm trying to land Z. The audience is..." Whatever comes out.

Then you say:

*"Take that mess and write me a 600-word draft of it. Do it in a confident professional voice. Don't worry about whether it's right — I'm going to use what's wrong about it to find what I actually want to say."*

You read what it produces. And here's the magic: you will instantly know what's wrong with it.

- Too clinical.
- Misses the point about Y.
- Sounds like a LinkedIn coach.
- The example is generic.
- This sentence is what I actually mean.

The wrongness is information. It tells you what your real voice wants to do that the average voice doesn't. The bits that survive your read are the bits that were already true. The bits that don't survive are signposts pointing at your real argument.

Often you write the actual piece in twenty minutes after this, because the throwaway draft has cleared the fog. The fastest way to find your voice is to read a draft that doesn't have it, and notice the gap.

## **A note on what makes these techniques work**

Both of these depend on something specific: you treating the model as a tool rather than an oracle.

If you ask AI what should I write about, you get bland answers because there's no specific information to work from. If you ask it here is my actual mess — help me see it more clearly, you get something useful, because you're providing the raw material only you can provide.

The pattern is consistent across every good AI use case. You bring the specific. The model brings the structure. Your context, your history, your real voice, your actual problem. The model holds it long enough for you to see it, and reflects something back you couldn't generate alone.

That's authorship. You staying central to the work, with a tool that makes your thinking visible to yourself.

## SECTION FOUR

## What Just Became Possible

---

There's a category of work that, until very recently, required either money or a technical team. A landing page that looks professional. A working prototype of a product idea. A simple app that does one thing well. A bot that answers questions on your behalf. A scheduling tool customised to your business. An internal dashboard.

Five years ago, building any of these meant either learning to code yourself — months of investment — or hiring someone, typically £5,000 to £50,000 depending on complexity. The cost wasn't just money. It was that most ideas died at the gate, because the friction of building them was higher than the conviction needed to start.

That gate has, quietly, been removed.

You can now describe what you want, in plain English, to a tool like Replit, Lovable, Bolt, or Claude Code, and have it build a working version. Not a sketch. Not a wireframe. A live, deployable thing. You can test it with real people. Refine it. Ship it.

This is genuinely new. And almost nobody you know is using it yet.

### What this looks like in practice

Imagine you've had an idea for a service. Maybe a tool for women returning from maternity leave to assess their options. Maybe a simple coaching prompt-of-the-day for your existing clients. Maybe a way for friends to share book recommendations that doesn't involve another social network.

The old path: write a business plan, save up, find a developer, brief them, wait three months, get something close to what you wanted, run out of money before you've tested it with real users. Most ideas die here.

The new path: open Lovable on a Saturday morning. Describe what you want. "I want a simple website where someone can answer five questions about their career situation and get a personalised summary back. I want it to look clean and modern. I want a way for them to share their email if they want to hear more from me." The

tool builds it. By Saturday afternoon you have a working version. Sunday you refine. Monday you share the link with five trusted people and ask what they think. By the end of the week you know whether the idea has any pull.

The whole process costs you a weekend and roughly nothing. You haven't quit your job. You haven't hired anyone. You haven't written a business plan. You've just made the thing exist long enough to test whether it should exist.

## What you can build right now without a developer

- A landing page that looks professionally designed.
- A working version of a small app or tool.
- A waitlist for an idea you want to test.
- A bot that answers FAQs about your business.
- A custom scheduling system tailored to how you actually work.
- An internal dashboard for tracking your own projects.
- A simple e-commerce setup for a digital product.
- A research tool that automates parts of your workflow.

None of these now require an engineer. Some of them you can describe in two sentences and have running before lunch.

## The shift this represents

The barrier to entry for testing an idea has collapsed. That's the headline. What it means in practice is that the bottleneck on creating something new is no longer money or technical skill. It is clarity. If you know what you want to make, you can probably make it. If you don't, no amount of tooling will help.

This puts the emphasis squarely back on the human work. What problem are you actually trying to solve? Who is it for? What does good look like? These are questions only you can answer. The technology executes. You direct.

For most of recent history, the people who could build things were a narrow class — engineers, designers, people with capital. Everyone

else had ideas they couldn't act on. That asymmetry is dissolving, faster than the public conversation has caught up with. The people who notice early have a window before the rest of the world does.

## A small caveat

You will not build a complex software business in a weekend. The tools are extraordinary at making small things quickly and surprisingly good at making medium things in a few weeks. They are not yet good at making large, complex, scalable systems without significant input. So calibrate your ambition: a landing page, a prototype, a working MVP — all genuinely doable. A full SaaS platform with thousands of users — still requires real engineering.

But the point is not to compete with funded startups. The point is to make your idea exist long enough to test it. That used to be the hard part. Now it's the easy part. The hard part has moved one square forward, to knowing what's worth building.

That is, again, an authorship question. Whose problem is this solving? Whose story is this telling? What do you actually want to make?

*The tools wait for your answer.*

## SECTION FIVE

## What AI Can't Do

---

*If you take only one thing from this guide, take this.*

The conversation about AI is overwhelmingly about what the technology can do. Each new model release, each new capability, each fresh anxiety about jobs and creativity and meaning, all of it driven by the same question: what can it do now that it couldn't do last month?

It's the wrong question.

The right question is the inverse. What can the technology not do, that I can? Because that, and not anything else, is where your value now lives.

The tools have absorbed an enormous amount of human output. They have read more books than any single person could in a hundred lifetimes. They can write competently in any genre, summarise any document, draft any kind of correspondence, produce passable versions of nearly any task that involves words or images.

### **What they cannot do is be you.**

They cannot have lived your particular life. They cannot have stood in your kitchen at three in the morning making a decision that shaped the next ten years. They cannot have lost what you have lost. They cannot carry your particular set of relationships — the friends who know your jokes, the family members you have complicated feelings about, the ex you parted from cleanly, the ex you didn't. They cannot have the body you walk around in, with its specific history of pleasure and damage. They cannot have your taste, which is the residue of everything you've ever paid attention to.

This is not a sentimental point. It is a strategic one.

The market is about to be flooded with AI-generated work. Average reports. Average essays. Average images. Average code. Average

everything, produced at zero marginal cost, in unlimited supply. Most of it will be technically competent and emotionally inert. The flood is already starting.

In a flooded market, what becomes valuable is what can't be flooded.

What can't be flooded:

- **Your specific lived experience.** Not the abstract version. The one that includes the wrong career move, the difficult parent, the friendship that ended badly, the year you didn't write, the conversation that changed everything. The specific texture of what you have actually been through. No model has been through it. No model can fake it.
- **Your taste.** What you find beautiful. What you find tedious. What you respond to in writing, in clothes, in food, in people, in argument. Taste is the accumulated weight of everything you have noticed, and noticing is irreducibly human work. A model can produce a thousand options. Only you can pick the right one.
- **Your context.** Who you know. Who trusts you. Who would take your call. What you understand about a particular industry, family, community, problem, that no document captures. Context is built over years and lives in your nervous system.
- **Your judgment under uncertainty.** When the data is incomplete, when the situation is novel, when the ethical question is real, the model can pattern-match but it cannot decide. Decision-making with skin in the game is yours.
- **Your voice.** Not the technical sense — the model can mimic style. The deeper sense. The thing that comes through when you write or speak that makes someone, somewhere, recognise you. That is a function of your specific history and cannot be generated.
- **Your presence.** What it is like to be in a room with you. What it is like to be loved by you. What it is like to be challenged by you. These are not transferable to a screen and they are not generatable from text. They are yours.

The standard advice in moments of technological disruption is learn the new tool. The standard advice is incomplete. The full advice is:

learn the new tool so well that you can use it to amplify the things only you can do.

Most people, panicking about AI, are running in the wrong direction. They are trying to become more like the machines, learning prompt engineering and tool stacks and workflow automation, in the hope that this will protect them. It won't, fully. The tools are commoditising. Everyone will have them. The differentiator cannot be access to the tool.

*The differentiator is what only you bring to the tool.*

That includes everything in your life that isn't on a page or a screen. The years of conversations that shaped how you think. The mistakes that sharpened your instincts. The grief that taught you what mattered. The relationships that are still, quietly, holding you up. The body you live in. The places you've been. The losses. The wins. The very specific way you happen to look at the world.

That is the asset. The technology is the leverage on the asset. Without you, the technology produces averages. With you, it produces something only your particular life could have made.

The reason this guide is called The AI Authorship Field Guide and not The AI Productivity Guide is this exact point. Productivity makes you faster at producing what could have been produced anyway. Authorship is what only you could have produced.

The era we are entering is going to ask, of every person who works with their mind: what is left of you, when the average has been automated?

The answer is: a great deal, if you have been paying attention to your own life.

This guide ends here, with that question, because everything else — the tools, the prompts, the techniques — is downstream of it.

*The technology is genuinely remarkable.  
But it is not the story.*

**You are.**

## SECTION SIX

## The 30-Day Starter Plan

---

The point of this guide isn't insight. The point is action that compounds.

Reading about AI changes nothing. Using it for thirty days, on small specific tasks that connect to your actual life, changes the way you think about it forever. By the end of the month you will not need any further guides, because you will have built your own working relationship with the tools.

What follows is a practical thirty-day plan. Each day takes between ten minutes and an hour. Most are closer to fifteen minutes. The plan assumes you have a single thinking partner set up — Claude or ChatGPT, free or paid, voice or text — and nothing else.

You can do these in order or skip around. The order is built so each one quietly prepares you for the next, but the plan is yours. Move through it the way you actually live.

### Week One: Getting comfortable

**Day 1 — Introduce yourself** — Open the model. Tell it who you are, what you do, what you're working on right now, and what you'd like help with over the coming month. Don't ask it to do anything specific yet. Notice how it responds. Notice what it gets and what it misses.

*Five minutes.*

**Day 2 — The voice-note dump** — Talk to the model for three minutes about something you've been thinking about. Anything. A decision you're weighing, a frustration, a project you can't quite start. End by asking: what patterns or themes are you noticing in what I just said? Read its response.

*Ten minutes.*

**Day 3 — The throwaway draft** — Pick something you need to write this week. An email, a post, a section of a document. Ask the model to write a draft, deliberately confident and slightly wrong. Read it. Notice what's wrong. Write your real version using the wrongness as a guide.

*Twenty minutes.*

**Day 4 — The pattern-spotting test** — Paste or read three recent messages you've sent. Ask: what do you notice about how I write to people? What might I be doing without realising? Sit with what comes back.

*Fifteen minutes.*

**Day 5 — The boring task** — Find something genuinely tedious that needs doing. Restructuring a long document. Summarising several articles. Generating a list of options. Hand it to the model. Notice how much faster it goes.

*Variable — whatever the task takes.*

**Day 6 — Reflect** — Look back at the week. Which of the days produced something useful? Which felt like a waste? Write down two sentences about what you've noticed about working this way.

*Ten minutes.*

**Day 7 — Rest, or build** — Take the day off, or use it to try one thing that's been on your mind that you'd never normally have time for.

*Free.*

## Week Two: Going deeper

**Day 8 — The future-self letter** — Ask the model to imagine you five years older, having made the bolder version of every choice you're currently weighing. Ask it to write you a one-paragraph letter from her about what she wishes you knew now. Read it twice.

*Fifteen minutes.*

**Day 9 — The honest critic** — Take a piece of work you've been pleased with — a recent post, a piece of writing, a presentation. Ask the model: what's wrong with this that I'm not seeing? Be specific and don't soften it.

*Twenty minutes.*

**Day 10 — The conversation rehearsal** — Pick a difficult conversation you need to have. Ask the model to play the other person and respond as that person would. Practise saying what you want to say. Refine.

*Twenty-five minutes.*

**Day 11 — The audit** — Give the model a description of how you currently spend your time at work. Ask: based on what I've told

you, what am I doing that I probably shouldn't be? What's missing that I probably should be doing?

*Twenty minutes.*

**Day 12 — The concept stretch** — Pick something you're vaguely interested in but haven't explored. Ask the model to teach you the basics in fifteen minutes, the way a kind expert would. Ask follow-up questions.

*Twenty minutes.*

**Day 13 — The creative experiment** — Make something for fun. A poem. A short story opener. A description of a place that doesn't exist. A pitch for a project you'll never do. The point is to write with the model as a sparring partner, not a productivity tool.

*Thirty minutes.*

**Day 14 — Reflect** — Two weeks in. What's surprised you? What's not landed?

*Ten minutes.*

## Week Three: Building

**Day 15 — The map of an idea** — Take an idea you've been carrying. A possible business, a project, a piece of writing. Ask the model to help you map it — what would it require, who would it serve, what's the smallest version of it?

*Thirty minutes.*

**Day 16 — The landing page** — Open Lovable, Bolt, or any of the no-code AI builders. Describe a simple landing page for the idea from yesterday. Watch it build one. Don't ship it. Just see it exist.

*Forty-five minutes.*

**Day 17 — The first draft of a manifesto** — What do you actually believe about the work you do? Talk to the model about it. Get it to draft a one-page manifesto. Edit it toward what you actually mean.

*Forty minutes.*

**Day 18 — The hard email** — Write the email you've been putting off. Use the model to draft three different versions in different registers. Pick the one closest to right and refine.

*Twenty minutes.*

**Day 19 — The reading list** — Ask the model to recommend five books you might love based on what you've already told it about yourself. Notice which ones spark something. Order one.

*Fifteen minutes.*

**Day 20 — The thinking partner** — Pick a real decision you're weighing. Have a thirty-minute conversation with the model about it — genuinely, not performatively. Don't ask it to decide. Ask it to help you see your own thinking.

*Thirty minutes.*

**Day 21 — Reflect, integrate, rest** —

*Ten minutes.*

## Week Four: Making it yours

**Day 22 — The script-spotting deep dive** — Show the model six pieces of correspondence that span the last year. Ask: what scripts are running through these? What am I consistently doing? Where is my voice and where am I performing?

*Thirty minutes.*

**Day 23 — The voice retrieval** — Find a piece of writing you did years ago that felt completely yours. Show it to the model alongside something recent. Ask: what's different? What's been lost? How do I get the first one back?

*Thirty minutes.*

**Day 24 — The small experiment** — Try something new with the tools — a use case you haven't explored, a different model, a feature you've been avoiding. Stretch the boundary.

*Variable.*

**Day 25 — The future map** — Imagine your work in three years. What does it look like? What are you known for? What would have had to happen between now and then? Have the model help you reverse-engineer the path.

*Forty minutes.*

**Day 26 — The contribution question** — Ask the model: based on everything I've shared with you over the past month, what do you think I should be contributing that I'm not yet? See what it says. Consider it.

*Twenty minutes.*

**Day 27 — The plan** — Take what's surfaced over the month and turn it into a one-page plan for the next three months. The model can help you structure it, but the content is yours.

*Forty minutes.*

**Day 28 — The handover** — Identify one task in your weekly routine that the model could meaningfully help with going forward.

Set up a system — a saved prompt, a regular workflow — so it becomes automatic.

*Thirty minutes.*

**Day 29 — The honest review** — What did this month change? What stayed the same? What was hype and what was real? Write three honest sentences.

*Ten minutes.*

**Day 30 — The first thing** — Decide one thing you're going to do, in the next week, that this month has shown you. Just one. Make it concrete. Begin.

*As long as it takes.*

## After day 30

The plan ends here, but the practice doesn't.

By now you will have a working relationship with the technology. You will know what it's good for and what it isn't. You will have your own preferred prompts, your own rhythms, your own sense of when to reach for the model and when not to.

That is the goal. Not to become an AI expert. Not to know every tool. Not to keep up.

The goal is to have integrated this set of capabilities into your life in a way that frees up more of your specific human attention for the things only you can do.

If you got this far, you have done that.

# The pen is yours.

## APPENDIX A

## The Toolbox

---

A working list of the tools worth knowing. You don't need all of them. You probably need three. This is the map so you can pick.

### Thinking partners (large language models)

The conversational tools you can talk or type to. General-purpose, do most of the work in this guide.

**Claude** (*claude.ai*) — Anthropic's model. My personal preference for nuanced thinking, longer writing, and conversations that need actual subtlety. Free tier available; paid is around £15/month for heavier use.

**ChatGPT** (*chatgpt.com*) — OpenAI's model. The most widely known. Strong all-rounder, particularly good at structured tasks and quick everyday questions. Free tier; paid (Plus) is £16/month.

**Google Gemini** (*gemini.google.com*) — Google's model. Integrated with Google Workspace if you live in Gmail and Docs. Good for research-style queries because it can search the web in real time.

**Meta AI** (*meta.ai*) — Built into WhatsApp and Instagram. Convenient if you already live there. Less powerful than the above but increasingly capable.

**Perplexity** (*perplexity.ai*) — Specialised for research. Cites its sources, which the others mostly don't. Useful when you need to know where information came from.

### Voice and audio

**ChatGPT voice mode** (*chatgpt.com*) — Built into the ChatGPT app. Talk to the model like a person. Genuinely changes how you use it.

**Claude voice mode** (*claude.ai*) — Also built into the Claude app. Same idea.

**ElevenLabs** (*elevenlabs.io*) — Generates extraordinarily realistic voiceovers from text. Useful for podcasting, narration, or creating audio versions of your writing.

**Otter** (*otter.ai*) — Transcribes meetings, voice notes, anything spoken. Generates summaries. Indispensable if you take a lot of calls.

## Image-makers

**Midjourney** (*midjourney.com*) — The benchmark for high-quality image generation. Subscription-based. Discord-native interface, slight learning curve, but the output is the best in the business for most aesthetic work.

**DALL-E** (*chatgpt.com*) — Built into ChatGPT. Easy access, decent quality, no separate subscription needed if you're already on ChatGPT.

**Imagen** (*gemini.google.com*) — Google's image generator, accessible inside Gemini.

**Nano Banana** (*nanobanana.ai*) — Newer, fast, free tier available, increasingly competitive with the bigger names.

**Stable Diffusion** (*stability.ai*) — Open-source, can be run locally, popular with people who want full control. Steeper learning curve.

**Recraft** (*recraft.ai*) — Specialised for design work, brand assets, graphics. Worth knowing if you need consistent visual identity.

## Builders (no-code AI development)

The tools that let you describe software and have it built. The most under-appreciated category in this guide.

**Lovable** (*lovable.dev*) — Describe a web app in plain English, watch it build. The fastest way to go from idea to working prototype.

**Bolt** (*bolt.new*) — Similar to Lovable. Sometimes one will produce better results than the other for a given idea, so it's worth trying both.

**Replit** (*replit.com*) — Long-established. Has its own AI agent (Replit Agent) that can build, debug, and deploy applications.

**v0** (*v0.dev*) — Vercel's tool. Particularly good at generating beautiful front-end interfaces.

**Cursor** (*cursor.com*) — A code editor that uses AI throughout. For people who want to move from no-code to actual coding gradually.

**Claude Code** (*claude.com/code*) — Anthropic's coding agent. Runs in your terminal. Builds and edits code on your machine.

## Agents (tools that take actions for you)

**Cowork** (*cowork.ai*) — Anthropic's desktop agent. Manages files and does tasks across your computer.

**Claude in Chrome** (*claude.ai*) — A browser extension that lets Claude navigate websites and take actions for you.

**Manus** (*manus.im*) — Autonomous agent that can plan and execute multi-step tasks.

**OpenAI Operator** (*openai.com*) — OpenAI's browser-based agent. Books things, fills forms, executes web tasks.

## Connectors and automation

**Zapier** (*zapier.com*) — Connects different apps so they trigger each other. Increasingly AI-augmented.

**Make** (*make.com*) — Similar to Zapier, more visual, often cheaper for complex workflows.

**n8n** (*n8n.io*) — Open-source automation. More technical but powerful for serious workflows.

## Specialist tools worth knowing

**Notebook LM** (*notebooklm.google.com*) — Upload your own documents and ask questions across them. Generates audio summaries that sound like a podcast about your own material.

**Granola** (*granola.ai*) — Note-taker for meetings that produces genuinely useful summaries.

**Pi** (*pi.ai*) — A more emotionally fluent conversational AI. Worth experimenting with for thinking-out-loud sessions.

**Suno** (*suno.com*) — Generates music from text descriptions. Useful for soundtracks, demos, fun.

**Runway** (*runway.ml*) — Video generation. Edit, extend, transform video clips with AI.

**HeyGen** (*heygen.com*) — Generates video of avatars speaking from text. Useful for course content and explainers.

## Where to start

If you've never used any of this:

- **Start with one thinking partner — Claude or ChatGPT.** Use it daily for a week. Notice what it's good at.

- **Add voice mode when you're comfortable typing.** It changes the relationship.
- **Try one builder (Lovable or Bolt).** On a Saturday when you have an idea you want to test.
- **Add one image tool.** When you have a creative reason to.

Don't try to learn everything. Tools are not the point. Use is the point.

## APPENDIX B

## Prompts To Steal

---

A working library of prompts you can copy directly into any thinking partner. Each one is designed to do a specific job. Use them as written, or adapt them to your situation.

The format below is the prompt itself in italics, followed by a one-line note on when to use it.

### Prompts for thinking

*"I'm going to talk through something I've been turning over. Don't try to fix it. Listen, then tell me what you notice about what I'm actually saying underneath the surface, what I seem to be avoiding, and what I'm more sure of than I'm letting myself admit."*

For when you have a tangled thought you can't quite see clearly.

*"I'm weighing a decision. I'm going to describe both options. Don't tell me what to do. Help me notice what each option is asking me to give up, and what I'm telling myself about each one that might not be true."*

For decisions where you're already leaning one way but won't admit it.

*"I keep coming back to this thought: [thought]. Help me work out whether it's a useful insight or a familiar loop. What questions could I ask myself to tell the difference?"*

For thoughts that feel important but might just be noise.

### Prompts for self-noticing

*"Here are three messages I've sent in the last week. Don't fix them. Tell me what patterns you notice in how I write to people — what I tend to do, what I might be avoiding, what I'm consistently softening."*

For spotting the scripts running through your communication.

*"Here is how I described my work to someone I just met. What assumptions am I making about what's important? What might I be hiding or downplaying? What am I emphasising more than I should?"*

For noticing how you're presenting yourself to the world.

*"Read this back to me but as if you were a kind, sharp friend who notices everything. What would she say about what I just wrote?"*

For when you need an outside eye on your own work and don't have one available.

## Prompts for writing

*"Take this messy thinking and write me a 600-word draft in a confident professional voice. Don't worry about whether it lands. I'm going to use what's wrong about it to find what I actually want to say."*

The throwaway-draft technique. Ironically the most useful prompt in the entire guide.

*"Write three different versions of this email in three different registers — warm, professional, and brief-and-firm. Show me how the same content lands differently."*

For difficult correspondence where tone is doing most of the work.

*"I want to write something about [topic]. The argument I'm trying to make is roughly [argument]. The audience is [audience]. Suggest five different angles I could take, and tell me what each one would emphasise."*

For when you know what to write about but not how to enter it.

*"Read this piece I've written. What's wrong with it that I'm not seeing? Be specific and don't soften it. I'd rather hear it now than after I've published it."*

For getting honest feedback on your own work without an editor.

## Prompts for difficult conversations

*"I need to have a difficult conversation with someone. I'm going to tell you who they are, what I want to say, and what I'm worried about. Then I want you to play the other person*

*and respond as they would. We'll do this several times so I can practise."*

For rehearsing conversations before they matter.

*"I've drafted this message but I'm not sure it lands the way I want. Read it. Tell me what the recipient is most likely to hear, and what they're most likely to misread. Suggest small changes that would close the gap between intent and reception."*

For tightening high-stakes correspondence.

*"I had a conversation that didn't go well. Here's roughly what was said. Help me work out what was actually happening — what they were trying to say, what I was trying to say, and where the misunderstanding sat."*

For making sense of a conversation in retrospect.

## Prompts for building and clarity

*"I have an idea. Let me describe it. Help me find the smallest version of it that I could test in a weekend. Don't give me a business plan. Give me the one experiment that would tell me whether to keep going."*

For ideas that need a reality check before they become projects.

*"I'm going to describe a piece of work I'm trying to make. Help me get clear on three things: what problem it actually solves, who specifically it's for, and what good looks like once it exists."*

For the first hour of any new project.

*"Here's what I'm planning to do over the next three months. Read it back to me and tell me what you think will actually happen versus what I'm hoping will happen. Where am I being unrealistic? Where am I being too cautious?"*

For sanity-checking your own plans.

## Prompts for the bigger picture

*"Imagine me five years older, having made the bolder choice on every decision I'm currently weighing. Write me a one-paragraph note from her about what she wishes I knew now."*

The future-self letter. Quietly powerful.

*"Based on everything you know about me from our conversations, what do you think I should be contributing that I'm not yet? Be specific."*

For the question we don't usually ask ourselves.

*"What's a story I keep telling about myself that might not be entirely true? Don't be cruel. Be accurate."*

For when you're ready to look at one of your own scripts.

*"What would it look like if I trusted my own judgment more? Use the things you know about me to be specific about which areas I'm currently outsourcing to other people's opinions."*

For taking back authorship one decision at a time.

## A note on using these

These prompts are starting points, not scripts. The model will respond differently each time. You can refine, push back, ask for more, ask for less. The conversation is the work.

The single most important habit is giving the model the specific raw material only you have. Real examples. Actual messages. Honest descriptions. The model will produce average results from average input. Specific input from your real life produces something only your life could have generated.

Treat the model as a sharp colleague who has limitless time but no context until you provide it. Then provide it generously. The depth of what you get back is a direct function of the honesty of what you put in.

## About the Author

---

Rebecca Holder is a writer, founder, and former senior corporate operator with twenty-five years across Accenture, the BBC, Barclays, and BP. Her first job was on Absolutely Fabulous. She went on to help launch BBC iPlayer and Barclays Faster Payments, so when someone finally repays a debt to your account in seconds, she'll happily take credit.

She is the founder of Real Humankind and writes the Substack *Whose Story Is It Anyway?* on storytelling, identity, and AI. She has four books almost ready to publish, has reversed two chronic conditions, has three children and a dog, and is currently rebuilding her life from the inside out, in public, in real time.

She works with senior professionals and organisations on the human side of AI integration, identity rewrites, and what it means to be exceptionally human in the age of machines that can generate everything except the truth of who we are.

### Stay in touch

**Substack:** [substack.com/@rebeccaholder](https://substack.com/@rebeccaholder)

**LinkedIn:** [linkedin.com/in/rebecca-holder-5a5273115](https://linkedin.com/in/rebecca-holder-5a5273115)

**Website:** [realhumankind.com](https://realhumankind.com)

**Email:** [rebecca@realhumankind.com](mailto:rebecca@realhumankind.com)

*Whose story is it anyway?*

**Yours.**

© 2026 Rebecca Holder · Real Humankind

*This guide is free to share. Please credit and link back when you do.*