A KELVINSTONE FIELD GUIDE

The Business Owner's *Introduction* to AI.

A plain-English guide to what it is, how it works, and what it actually means for your business. Written for people with ten other things on their plate.

PUBLISHED

Spring 2026 · Glasgow, UK

What's *inside*.

Eight short sections plus a glossary, designed to be read in one sitting or in small chunks over a week. No jargon, no hype, no selling at you.

01

Introduction

Why this guide exists, who it's for, and what you'll get out of reading it. Honest framing — no hype, no selling, no ten-hour YouTube rabbit holes.

WHY THIS GUIDE EXISTS

Most AI writing for business owners is *terrifying*, overhyped, or trying to sell you something.

This one tries not to be any of those things. It's a free guide, written for people running a business who've heard about AI for years without really knowing what to do about it.

You don't need to be technical. You don't need to have used any of it. By the end, you'll understand what AI actually is, how it works under the hood — enough to make sensible decisions about it — what it's genuinely useful for, what it isn't, and where it's probably heading.

Just what you need to know about how to start understanding and taking advantage of this technology.

Who this is for

- Business owners who haven't really used AI yet, or have only dabbled with ChatGPT a few times.
- Anyone who feels like they're falling behind on this but doesn't know where to start.
- People who want a straight answer instead of ten hours of YouTube videos full of guys shouting at a camera.

If you're already building AI systems or running a tech company, this probably isn't for you. It's written for the person who's got a business to run and twenty other things on their plate.

Who we are

Kelvinstone AI is a small agency that builds AI tools and automations for small businesses. That's the short version.

The longer version: big companies are going to use AI to pull even further ahead of everyone else. That's already happening. Kelvinstone exists so small businesses don't get left behind in that shift. We build practical stuff for the kind of businesses that don't have an IT department or a tech budget, and we explain what we're doing in language that makes sense.

That's why this guide exists too. Giving away good information for free is part of the point. If you read it and decide you want help implementing any of it, get in touch. If you read it and sort things out yourself, that's also a win.

What you'll get out of this

By the end of the guide you'll have:

- A clear mental model of what AI actually is — and what it isn't.
- Enough understanding of how the technology works to make sensible decisions about it.
- A sense of what's being automated in small businesses right now, and what isn't worth bothering with.
- Honest information about the limitations and risks — the stuff the hype merchants leave out.
- A view on where this is heading and what it means for you.

The guide is designed to be read in one sitting, or in short chunks over a week. It's around fifty pages. There's a second guide that covers the practical side — how to actually use AI day-to-day, prompts that work, and how to spot automation opportunities in your own business. You can grab that one too once you've finished this.

Let's get into it.

02

The Shift That's Happening

AI has moved from research project to mainstream tool in about three years. The cost is falling, the tools are simple, and your competitors are already using it.

54%

of UK small and medium businesses now use AI.

Two years ago, it was 25%.

SOURCE: BRITISH CHAMBERS OF COMMERCE, MARCH 2026

WHERE WE ARE NOW

Three years ago, AI in your business meant a developer, a custom build, and a *decent budget*. Most small businesses ignored it. That was the right call.

That's no longer the case.

The tools are now cheap or free. The interfaces are simple enough that anyone can use them. And businesses that have started implementing AI are already seeing real, measurable returns — not speculative future benefits, but time and money saved this quarter.

What follows is what the UK data actually says, what small businesses are using AI for, and what it means if you're not using it yet.

The UK picture

The numbers vary depending on who's counting and how, but the direction is the same across every credible source. A few worth knowing.

In March 2026, the British Chambers of Commerce reported that 54% of UK small and medium-sized businesses are now actively using AI. Two years earlier, that figure was 25%. Adoption has more than doubled, and it's still climbing.

The shift is showing up in how owners talk about it too. In a UK government survey, business owners described AI as *"imperative to survival"* and *"something you have to use to stay competitive."* That's a notable change in tone from a few years ago, when most small business owners would have shrugged at the question.

> **WHAT THIS MEANS**
>
> The headline number hides something more important — and it's the bit that should actually concern you if you run a small business. The next page breaks it down.

The two-speed race

The Bennett School of Public Policy at the University of Cambridge, analysing Office for National Statistics data, found that between 2023 and 2025, large UK firms (250+ employees) nearly doubled their AI adoption to 44%. Small firms (fewer than 50 employees) gained far less ground, reaching just 26%. Their conclusion was blunt: AI uptake in the UK looks less like a rising tide and more like a two-speed race, with larger firms pulling further ahead.

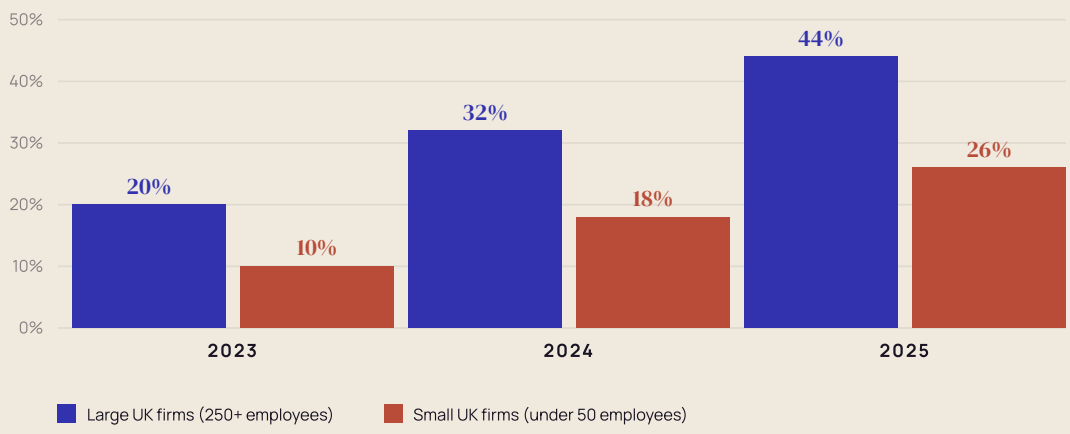

In the UK, large firms are pulling ahead faster than small ones are catching up.

The reason is structural, not technical. Bigger companies have the budget, the staff, and the time to figure this stuff out. Smaller businesses are running flat out keeping the lights on, and "spend a few weeks researching AI" doesn't make the priority list.

What AI is actually being used for

Most of the numbers above are about adoption — whether businesses are using AI at all. The more useful question is what they're using it for. Among UK businesses that *are* using AI, government research (DSIT, 2025) found the most common uses for small and medium firms break down like this:

Creative and content creation	77%
Administrative tasks (scheduling, invoicing, documentation)	70%
Data analysis	56%
Customer service (chatbots, enquiry handling)	46%

The top use cases aren't exotic. They're the boring, repetitive parts of running a business — admin, marketing, customer service, enquiry handling. The kind of stuff that eats hours every week and that nobody actually enjoys doing.

And the time saved is real. A 2025 UK study by Foxit found workers using AI saved **5.2 hours per week** on average — though close to four of those hours went back into checking and correcting the output. The net gain is genuine, but smaller than the headlines suggest. Treat AI as a strong drafter, not a finished worker.

The honest bit most guides leave out

Here's where we have to be straight with you, because a lot of the AI content online glosses over this.

Not everyone who adopts AI is winning with it.

A 2026 PwC study of over 1,200 senior executives found that nearly three-quarters (74%) of AI's economic value is captured by just 20% of organisations. The majority of businesses using AI are seeing modest gains at best — small efficiency improvements, a bit of time saved — while a smaller group is pulling seriously ahead.

The difference isn't the tools. Everyone has access to the same tools. The difference is how they're being used.

The businesses getting the biggest returns are the ones who:

- Picked one specific problem to solve first, instead of trying to automate everything at once.
- Built workflows that actually remove work, rather than just adding new tools to manage.
- Understood the technology well enough to know where it fits and where it doesn't.
- Treated AI as something to integrate into how they already work, not a separate thing they bolt on.

None of that requires being technical. It requires being clear-headed about what you're trying to achieve, and having enough understanding of the technology to make sensible calls. Which is most of the point of this guide.

Where this leaves you

If you're reading this, you're probably somewhere in one of three places.

GROUP A · HAVEN'T STARTED

You haven't used AI at all yet. You're not behind yet, but you will be if you wait another year. The gap is going to keep widening, and the longer you leave it, the harder it is to catch up.

GROUP B · DABBLED

You've dabbled. Used ChatGPT a few times, tried a tool or two. This is most people. Dabbling is fine as a starting point, but on its own it doesn't move the needle. The returns come from picking a specific problem and actually implementing a solution.

GROUP C · IMPLEMENTED

You've implemented something and it's working. Good. The next question is where else it could apply in your business — and that gets easier the more you understand the underlying technology, which is what the rest of this guide covers.

Whichever group you're in, the rest of this guide is designed to give you enough understanding to make sensible decisions. Not to turn you into an AI engineer. Just to stop AI from being a black box you can't evaluate. Let's start with what it actually is.

03

What AI Actually Is

The mental model. AI vs. AGI, how it all fits together, why LLMs are just prediction machines, and the handful of concepts you need to stop treating this as a black box.

THE MENTAL MODEL

The term “AI” gets used so loosely it’s basically *meaningless*.

People use it to mean ChatGPT, the recommendation algorithm on TikTok, the autopilot on a Tesla, and the spam filter on Gmail. All of those are technically AI. None of them are the same thing.

By the end of this section you’ll understand:

- The difference between AI and AGI (and why it matters).
- How AI, machine learning, deep learning, and generative AI all fit together.
- What an LLM actually is and how it works.
- What tokens and context windows are.
- The difference between closed and open-source AI models.

Some of this is a bit dense. Take it in chunks if you need to. The good news is that once you’ve got the basic concepts, the rest of the guide is mostly about applying them.

AI vs. AGI

You'll hear two terms thrown around: AI and AGI. They're related but they mean different things.

AI (Artificial Intelligence) is what we have now. The broad field of building systems that can do things we'd normally need a human brain to do — recognising images, understanding language, making predictions, generating text. AI is everywhere already: your bank's fraud detection, your phone's voice assistant, the chatbot on a delivery company's website, the tool that auto-edits your photos.

The key thing about today's AI is that it's *narrow*. Each system is built for a specific task and can only do that task. ChatGPT is brilliant at language but couldn't drive a car. A self-driving system can navigate a motorway but can't write you an email. They're specialists, not generalists.

AGI (Artificial General Intelligence) is the theoretical version that doesn't exist yet — a system that can do any intellectual task a human can do, across any domain, without being specifically trained for it. The goal of the big labs is to build AGI, and they're basically in an arms race to get there first.

> WHAT THIS MEANS FOR YOU
>
> Whether AGI is months, decades, or never away depends on who you ask. For running a business, it doesn't matter yet — you'll be working with narrow AI for the foreseeable future. For the rest of this guide, when we say "AI" we mean what's available today.

How it all fits together

Four terms you'll hear constantly — AI, machine learning, deep learning, and generative AI — all sound similar but mean different things. They're nested inside each other.

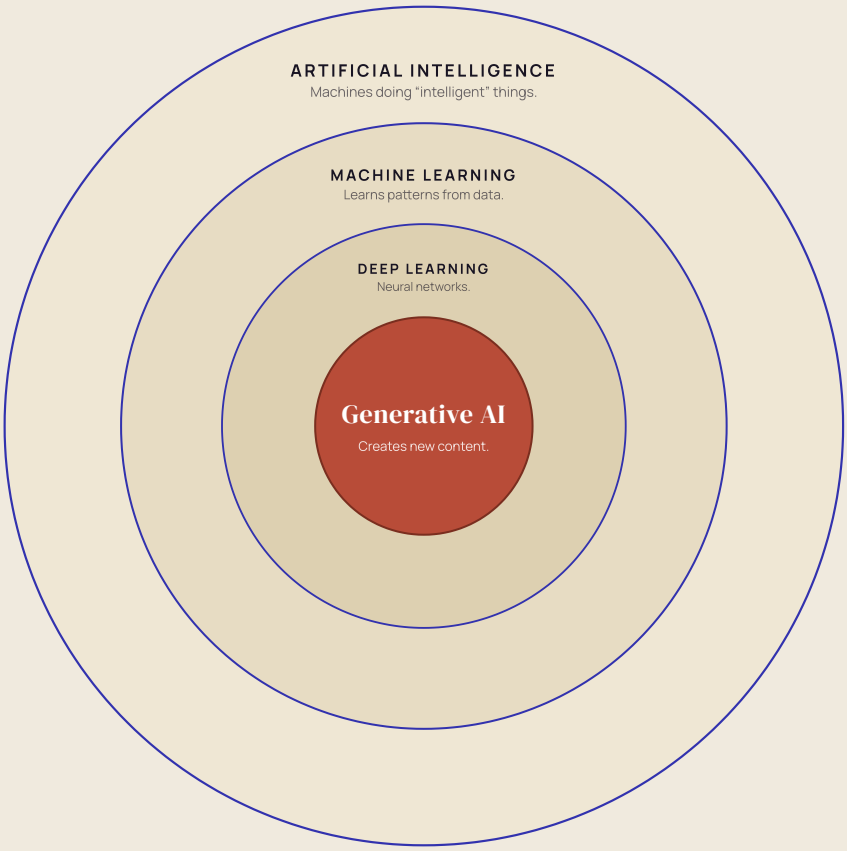

Four terms nested inside each other. When people today say "AI", they usually mean the innermost circle.

Machine Learning sits inside AI: instead of explicit rules, you give the system data and let it learn the patterns. Netflix recommendations. Fraud detection. Spam filters. **Deep Learning** sits inside that — ML using neural networks, which is where AI got genuinely impressive around 2012. **Generative AI** is the innermost ring: deep learning systems that don't just classify or predict, they *create*— text, images, video, audio, code. When most people today say "AI", this is usually what they mean.

THE SINGLE MOST IMPORTANT CONCEPT

LLMs are *prediction machines*. That's it. That's what they do.

When you type a question into ChatGPT, it doesn't "understand" your question the way a human does. It doesn't think about it. It predicts what words should come next, given the words that came before — one token at a time — until it's generated a complete response.

The closest everyday comparison is the autocomplete on your phone.

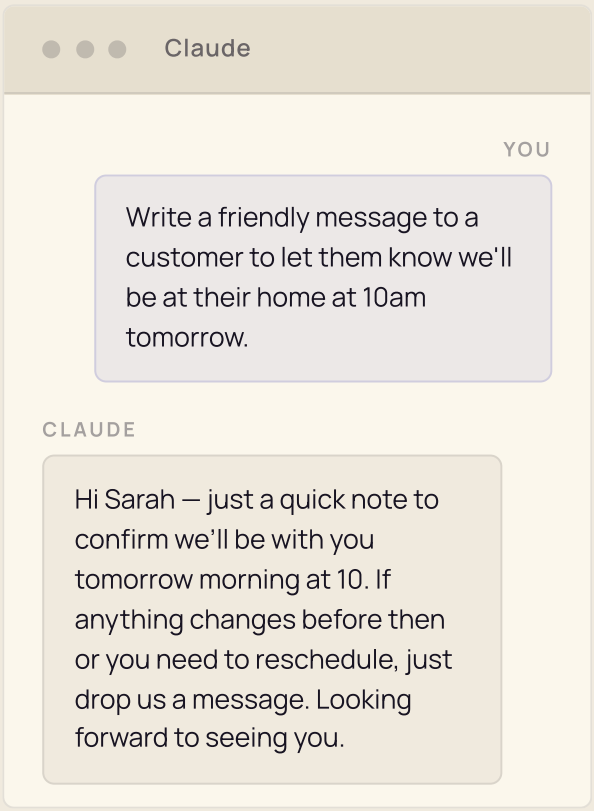

Same underlying idea. Predict what comes next. The difference is scale.

Why this matters practically

It sounds underwhelming when you put it like that. *"It's just predicting the next word? How is that doing all the impressive stuff I've seen?"* The answer: turns out, if you predict next words really well — based on having read most of the text on the internet — you can do a surprising amount of useful things. A coherent email. A summary of a long document. An explanation of a complex topic. Not because the AI understands any of it, but because it's seen enough examples.

Once you understand this, three things become obvious:

1 · THEY CAN BE WRONG, CONFIDENTLY

They predict what's *likely* to come next, not what's *true*. If they've seen lots of text saying something incorrect, they'll predict that incorrect thing back to you with total confidence. This is called "hallucination" — covered properly in Section 7.

2 · THEY CAN'T READ YOUR MIND

They only have what you give them in the prompt. The more context you provide, the better the predictions get. Vague prompts produce vague answers. This is why prompt quality matters so much — covered in Guide 2.

3 · THEY DON'T ACTUALLY KNOW THINGS

They've absorbed patterns from training data. There's no internal database of facts they're looking things up in. Which is why they sometimes get basic facts wrong — they're predicting plausible-sounding answers, not retrieving verified information.

Tokens, context windows, and training data

A few more concepts you'll bump into.

Tokens. LLMs don't read words the way you do. They break text up into smaller chunks. Common words might be one token; longer or unusual words get split into pieces.

"The plumber fixed the leaking tap quickly."

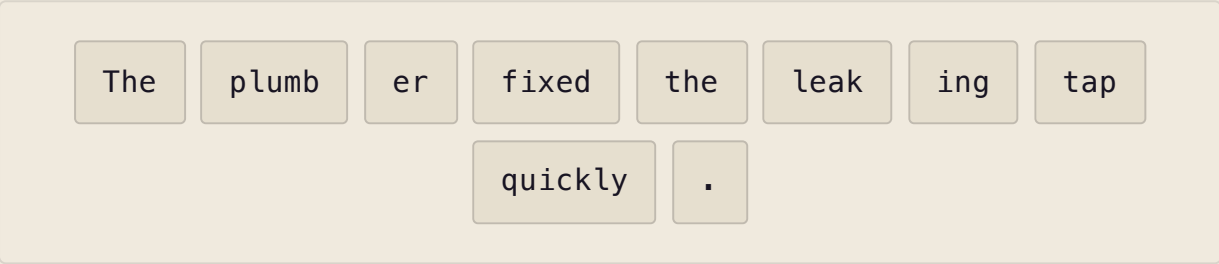

Ten tokens, not nine words. Tokens average about three-quarters of a word.

Context window. The maximum amount of text an AI can hold in its "working memory" at once — your prompt plus its response, plus any documents, plus the conversation so far. Claude can handle hundreds of thousands of tokens at once. ChatGPT's free tier is more limited. When you exceed the window, the AI starts forgetting the earliest parts — sometimes called *context rot* (Section 7).

Training data. What the model learned from: roughly the public internet, large numbers of books, Wikipedia, scientific papers, code repositories. Hundreds of billions of words. Models are frozen at a cutoff date, reflect the biases of what they were trained on, and know nothing about your business unless you tell them.

Closed vs. open-source models

One last distinction worth knowing about.

	CLOSED SOURCE	OPEN SOURCE
Examples	ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google), Grok (xAI)	Llama (Meta), DeepSeek, Mistral
Who controls it	The company that built it — weights, training, and interfaces all proprietary.	Anyone. Weights are released publicly; can be modified and redistributed.
Who can run it	Only the provider. You access it through their apps or API.	Anyone with the hardware. Runs on your own computer or servers.

For most small businesses this distinction doesn't matter much day-to-day — you'll use closed-source models like Claude or ChatGPT because they're the most capable and the easiest to use. It starts to matter when you care about **data privacy** (open-source models can be run fully in-house), **cost at scale** (closed-model API bills add up fast), or **customisation** (closed models can't be fine-tuned).

> SELF-HOSTING
>
> Running an AI model on your own computer or servers, rather than a cloud service. Possible with open-source models; overkill for most small businesses; worth knowing the option exists if you handle highly sensitive data (legal, medical, financial). Get someone technical involved if that applies.

Section recap

That was a lot. The short version of what to remember:

- **AI** is a broad term covering any attempt to make computers do "intelligent" things. **AGI** is the theoretical version that can do everything a human can. We're nowhere near AGI.
- **Machine learning, deep learning, and generative AI** are nested inside each other like Russian dolls. Generative AI is the bit causing all the recent fuss.
- **LLMs are prediction machines.** They predict what comes next based on patterns. They're not thinking, understanding, or knowing — they're predicting. This explains most of their strengths and most of their weaknesses.
- **Tokens** are how AI reads and writes text. **Context windows** are how much they can hold in working memory at once. Bigger is better.
- **Training data** is what they learned from. They're frozen at a point in time and don't know your business unless you tell them.
- **Closed vs. open source.** For most small businesses, you'll use closed source models like Claude or ChatGPT. Open source matters for privacy, cost at scale, and customisation.

That's the foundation. Next, we'll look at the actual tools available — what they're called, what they're for, and how to think about which ones to pay attention to.

04

The Current Tool Landscape

A field guide, not an encyclopedia. The three AI assistants worth knowing, what else you'll hear about, and why the principle matters more than the product.

A FIELD GUIDE, NOT A SPEC SHEET

There are dozens of AI tools out there. Most articles trying to cover them all end up as *exhausting feature comparisons* nobody reads.

This one won't do that.

By the end of this section you'll know:

- The three main AI assistants worth knowing about, and where each one is genuinely better.
- What the picture looks like for image and video generation.
- Why every software company is suddenly stuffing "AI" into their products — and why most of it isn't worth getting excited about.
- What "vibe coding" means, and how a separate category of developer tools has quietly changed who can build software.
- The most important principle for choosing tools: don't marry one.

Let's start with the assistants you'll actually use.

The three main AI assistants

If you've been paying attention, you've heard of ChatGPT. Maybe Claude and Gemini too. These are the three that matter for most small businesses.

The three you'll actually use.

Claude

ANTHROPIC

BEST FOR

Long documents, careful writing, coding, anything where accuracy matters.

WATCH OUT FOR

Smaller plugin ecosystem, no built-in image generation.

Free tier · £15/mo Pro

ChatGPT

OPENAI

BEST FOR

General use, brainstorming, quickest to learn, biggest plugin library.

WATCH OUT FOR

Occasionally drifts on long, multi-step prompts.

Free tier · £15/mo Plus

Gemini

GOOGLE

BEST FOR

Anyone living in Google Workspace. Very large documents, multimodal tasks.

WATCH OUT FOR

Less consistent quality round-to-round.

Free tier · £15/mo Pro

They all do roughly the same thing: type something, get a response. All three have free tiers that are genuinely useful, and paid tiers around £15 a month. Pricing is essentially identical. But they're not interchangeable — each has developed real strengths, and using the right one for the right job makes a noticeable difference.

Claude is the one we use most at Kelvinstone. Strongest at writing quality, long documents, and accuracy. Hallucinates less than the others. It's what we'd recommend if you're going to pick one primarily — though we're not married to it. If something better comes along, we'll switch without hesitation.

The other names you'll hear

Beyond the main three, there are a few others worth knowing about — mostly so you understand what people are talking about, not because you necessarily need to use them.

Grok is xAI's model (Elon Musk's AI company), integrated into X and pulling heavily from real-time data on the platform. If you're not heavily on X, there's not much reason to use it over the main three.

DeepSeek is a Chinese AI company that made waves in early 2025 by releasing open-source models competitive with the big Western players at much lower cost. It shifted the conversation about how expensive AI has to be — but for day-to-day small business use, you'll probably never touch it directly.

Llama is Meta's open-source model family, released freely. Unlikely you'll use it directly, but you may use products built on top of it.

Microsoft Copilot is baked into Office 365. Essentially ChatGPT under the hood, wrapped to work inside Word, Excel, Outlook, and Teams. If your business lives in Microsoft's apps, this is the obvious option.

— A PARALLEL TOOLSET, AND A NEW WAY OF WORKING

Developer tools — and the rise of "vibe coding"

Alongside the chat assistants, each of the big three has released a separate tool aimed squarely at people who write software: **Claude Code** from Anthropic, **Codex** from OpenAI, and **Gemini CLI**from Google. These don't live inside a chat window. They run inside a programmer's terminal or editor, and they can read, write, and change the code on a computer directly. They are how most serious AI automation actually gets built today.

Over the last year, one piece of language has broken out from the developer world into general use: **"vibe coding."** It describes what happens when someone without a traditional programming background uses one of these tools to build real, working software — describing the outcome they want in plain English, and letting the AI handle the syntax. It has quietly shifted who can build what. Internal tools, small apps, custom automations — the kind of thing that would once have needed a six-figure software project is now being put together by small teams in a matter of days.

At Kelvinstone, Claude Code is the tool we specialise in. We believe it's currently the most powerful of the three — strongest at following long, complex instructions, and the least likely to produce code that looks right but isn't. Codex and Gemini CLI are both serious alternatives, and the gap narrows every few months. As with the chat assistants, we're not loyal to the logo. We're loyal to what works.

For running your business day to day, you don't need to touch any of these yourself. Knowing they exist is enough. But it's worth understanding that when a small agency can now build something a ten-person development team would have quoted you six figures for a couple of years ago — this is how.

Image, video, and other media generation

The conversational tools are only half the story. There's a separate world of AI that generates images, video, audio, and other media — and it moves even faster than chat. The leaders change every few months.

For images: Nano Banana 2 (Google's current flagship, integrated into Gemini), Midjourney (long the favourite of designers, mostly through Discord), and Adobe Firefly (built into Photoshop and the rest of Adobe's suite). For most small business needs — social graphics, mock-ups, simple marketing assets — any of these does the job.

For video: Veo 3.1 (Google's leader, generates clips with synchronised audio), Runway and Kling (longer-form video with more creative control). Until very recently, OpenAI's Sora dominated the conversation — until OpenAI announced its shutdown in March 2026. More on that on the next page.

For other media: ElevenLabs for realistic voice synthesis, Descript for podcast editing, Suno for full song generation, HeyGen and Synthesia for avatar-based explainer videos.

You don't need to know all of these. They're listed so the names aren't unfamiliar when you bump into them. For most small businesses, the priority should be getting one chat assistant working well first.

The Sora story (and why you shouldn't marry a tool)

In February 2024, OpenAI announced Sora — an AI video generation model. The demos were stunning. Disney was reportedly planning a $1 billion investment in OpenAI partly off the back of it. Sora 1 launched publicly in December 2024 at $20/month; Sora 2 followed on 30 September 2025. By early 2026 the platform had over four million active users.

On 24 March 2026, OpenAI announced they were pulling the plug. The consumer Sora app shut down on 26 April 2026; the Sora API follows on 24 September 2026.

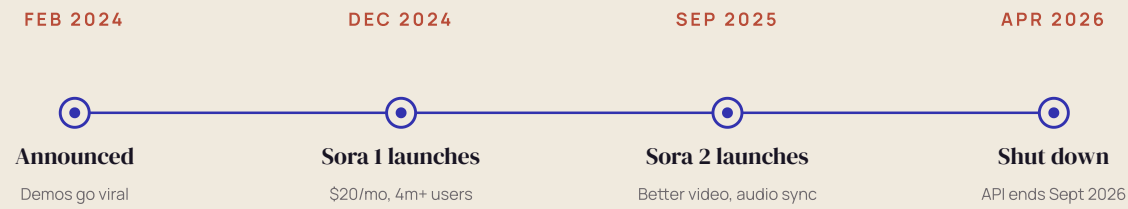

From breakthrough to deprecated in 26 months. This is why you don't build your business around one tool.

The reasons were essentially economic: Sora reportedly cost OpenAI $8–12 million a month in compute against $2 million in revenue. They decided the resources were better spent elsewhere. The tool millions of people had built workflows around was switched off with about a month's notice for the consumer app, and a six-month runway for anyone using the API.

This isn't a story about Sora being a bad product. It's a story about what happens when you build your business around a single tool from a single company. The lesson: **understand the principles, not just the tools.** The tools will keep changing. The principles won't.

— A NOTE ON THE PACE OF CHANGE

These tools will look different in six months. That's the point.

The AI field is moving faster than almost any technology shift of the last thirty years — and the rate of change is *accelerating*, not slowing down. Models released eighteen months ago have already been matched or beaten by free, open-source alternatives. Tools that dominated the conversation last year have been shut down, replaced, or quietly absorbed into something bigger. New flagship versions arrive every few weeks. What's best-in-class today is very unlikely to still be best-in-class by the end of the year.

That has to shape how you think about all of this. It's tempting to pick a tool, learn its quirks, wire your workflows around it, and quietly start treating it as "our AI." **Don't.** Every tool you use today is a current-best-option, not a permanent choice. Keep your processes portable. Stay interested in what's new without being seduced by every headline.

> ***The question is never "am I using the newest thing?" — it's "am I getting the outcome I care about?"***

That's the scoreboard that matters. Hours saved. Mistakes avoided. Work you couldn't take on before. Revenue that didn't exist last year. Focus on those, and the question of which logo sits at the top of the screen becomes a detail — not a strategy.

Why "AI" is suddenly in everything

Walk through any software company's marketing in 2026 and you'll struggle to find a product that doesn't claim to be "AI-powered". Photoshop has AI. Notion has AI. Your accounting software has AI. The app you book restaurant tables with probably has AI now too.

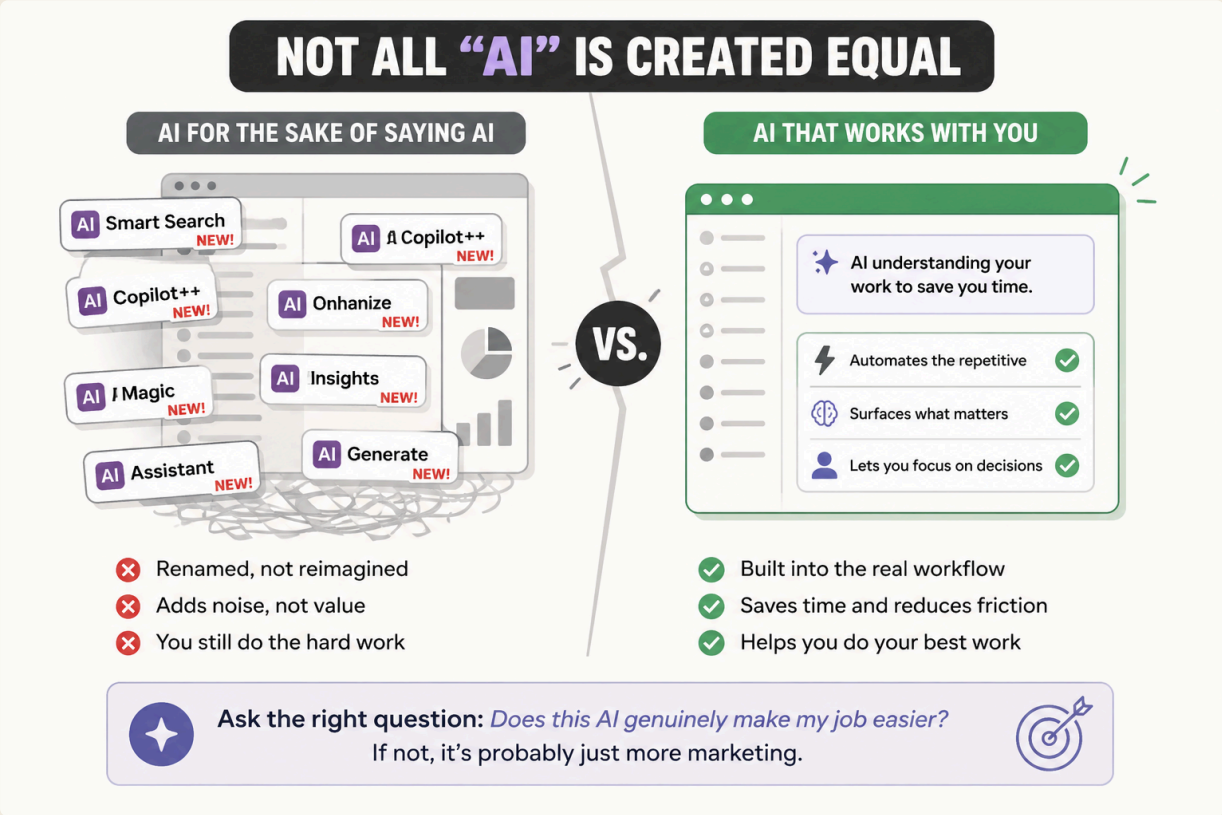

Most of it is marketing decoration. A company struggling to compete slaps "AI" on every feature, existing things get renamed ("Smart Search" becomes "AI-Powered Search"), a chatbot appears in the corner, and usually none of it actually makes the product better. The genuine exceptions — Adobe Firefly inside Photoshop, Microsoft Copilot inside Office — are tools where AI is added thoughtfully to work people already do.

When evaluating anything claiming AI features, ask yourself: *is this making my actual job easier, or is this just AI for the sake of saying you have AI?* If you can't answer clearly, it probably isn't worth paying extra for.

Section recap

- **Three main AI assistants** worth knowing: Claude (best for writing, accuracy, long docs), ChatGPT (most versatile, biggest ecosystem), Gemini (best if you're in Google Workspace). All around £15/month paid, all with usable free tiers.
- **Other names** you'll hear (Grok, DeepSeek, Llama, Microsoft Copilot) — useful to recognise, but for most small businesses the main three cover the ground.
- **Developer tools and "vibe coding"** (Claude Code, Codex, Gemini CLI) are a parallel category aimed at programming work. You won't use them directly, but they're how small agencies — Kelvinstone included — now build custom software that would have cost six figures a few years back. Claude Code is the one we specialise in and, for now, the most powerful of the three.
- **Image and video generation** are separate categories with their own leaders (Nano Banana 2, Midjourney, Veo 3.1). The space moves fast — what's best today won't be in six months.
- **The Sora shutdown** is one example, not a one-off. Companies pivot, products get killed, and your business shouldn't depend on any one of them.
- **The pace of change is accelerating.** Don't get hung up on any specific tool. Focus on the outcome — hours saved, mistakes avoided, work you couldn't take on before — and the logo at the top of the screen becomes a detail, not a strategy.
- **Be sceptical of "AI" badges** stuffed onto existing products. Some of it is genuinely useful (Adobe Firefly, Microsoft Copilot). Most is marketing noise that adds friction rather than removing it.

That's the landscape. Next, we cover something almost no other guide bothers to explain: the difference between AI automation and old-fashioned deterministic automation — and why understanding the difference is one of the most useful things a small business owner can do.

05

Deterministic vs. AI Automation

Not every problem needs AI. The difference between old-school automation and the AI kind — and why the most useful systems quietly combine the two.

— NOT EVERY PROBLEM NEEDS AI

Here's something almost nobody tells you: *not every problem needs AI.*

The AI hype machine has spent the last three years convincing everyone that AI is the answer to every business problem. It isn't.

A lot of what makes a small business inefficient can be solved with much simpler tools — automation that's been around for a decade and costs almost nothing to run. This section is about knowing when to reach for which, so you don't spend money using AI for a job a £10-a-month tool could do on its own.

By the end of this section you'll know:

- The difference between deterministic and AI automation.
- When to use which — and when to use both together.
- How every automation is structured, underneath the jargon.
- Why the most useful systems combine the two.

This is the section that separates people who can think clearly about automation from people who just throw "AI" at every problem and hope for the best.

Two different kinds of automation

Deterministic automation is the old school — and most small businesses already touch it without realising. The principle is simple: *if this happens, do that.* A form gets filled out, save the data to a spreadsheet. An invoice gets paid, send a thank-you email. No thinking, no interpretation; the same action every time. It's cheap, reliable, and predictable. The one thing it can't do is handle anything that needs interpretation — if the situation doesn't match the rules, it either does nothing or does the wrong thing.

AI automation is the newer category. It drops a large language model — Claude or ChatGPT, for example — into the workflow to handle the bits that need judgment. An enquiry arrives: is it a qualified lead or spam? About something you offer, or something unrelated? An LLM can read the message and make that call in a way deterministic rules simply can't. The trade-off is cost (a little per run), speed (seconds, not milliseconds), and consistency (the same input won't always give an identical output). Worth it for judgment work. Wasteful when the job is just moving data.

START HERE

Does this task need judgment, language understanding, or interpretation?

YES

Use AI

Drafting an email reply. Summarising a document. Deciding if an enquiry is qualified.

NO

Use deterministic automation

Copying a form submission to a spreadsheet. Sending a fixed reminder. Posting to Slack when a payment lands.

Both are valuable. Most businesses need a mix.

The right question isn't "deterministic or AI?" — it's *"which parts of this workflow need judgment?"* Use AI only there. Use deterministic automation for everything else.

The basic shape of any workflow

Whether a workflow uses AI or not, every automation has the same three-part structure.

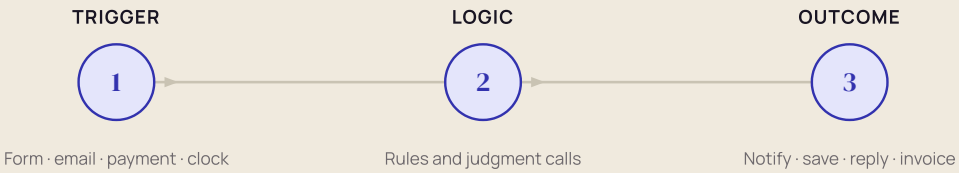

The trigger is what starts it. Something happens and the system kicks in. A form gets submitted. An email arrives. A payment lands. The clock hits 9am. A file appears in a folder. Without a trigger, nothing fires.

The logic is the middle bit — what the system does once it's running. Deterministic logic is rules (*if the email is from a known customer, file it under their name*). AI logic is judgment (*read the email and decide whether it needs an urgent reply*). Most workflows chain several logic steps together.

The outcome is the useful thing at the other end. A WhatsApp message to the owner. A calendar event created. An invoice drafted. A file saved to the right folder. That outcome is the whole point — everything before it is plumbing to make it happen without anyone having to do it by hand.

Once you see the pattern, you start spotting opportunities for automation everywhere. Any time you have a repetitive *"every time X happens, I have to do Y"* situation, you've got a candidate for a workflow.

A worked example

A small business gets enquiries through a contact form. The owner currently checks submissions by hand, decides which are worth following up, and writes each reply — 30 to 45 minutes a day, and new leads sometimes wait 24 hours for a response. Here's what an automated version might look like.

TRIGGER
Customer fills in the contact form
The submission kicks the whole workflow off.

DETERMINISTIC
Submission saved to a Google Sheet
Straight data copy. No interpretation needed.

AI-POWERED
AI reads the enquiry and decides if it's qualified
Real job, spam, or wrong fit? The model makes the call the way a human skim-reader would.

AI-POWERED
AI drafts a personalised reply
Uses the enquiry details and your tone of voice to write a suggested response.

DETERMINISTIC
Owner gets a WhatsApp with summary and draft
Plumbing — get the result in front of the human.

HUMAN
Owner reviews, edits, and sends
Anything touching the customer still passes through a person.

A real workflow that mixes deterministic, AI, and human steps. Each part does what it's best at.

Notice how AI is only used where judgment is required. The spreadsheet write, the notification, the trigger itself — all deterministic, all cheap and fast. The two AI steps are doing the work an LLM is actually good at: reading language and drafting a contextual reply. And anything that touches the customer still goes through the owner.

30 to 45 minutes a day becomes around five minutes reviewing drafts, and new leads get a response in minutes rather than hours — which makes a real difference to conversion.

The tools that let you build this

A year ago, this section would have been about no-code workflow builders — Zapier, Make, n8n. They still exist and are fine for simple jobs. For anything more ambitious, they've been overtaken by **Claude Code**.

zapier | make | n8n

Claude Code is a different kind of tool. You describe what you want a workflow to do in plain English, and it writes and runs the custom code that makes it happen — no platform integrations to wait for, no rigid drag-and-drop canvas. If an API exists, Claude Code can use it. That makes it the most versatile option in this space, and what we build almost everything on at Kelvinstone now.

The trade-off is that someone capable still needs to sit behind it, knowing what to ask for and how to maintain what gets built. That's most of what we do. The principles in this guide apply whether you build the systems yourself, hire us, or hire someone else.

WHAT WE BUILD ON TODAY

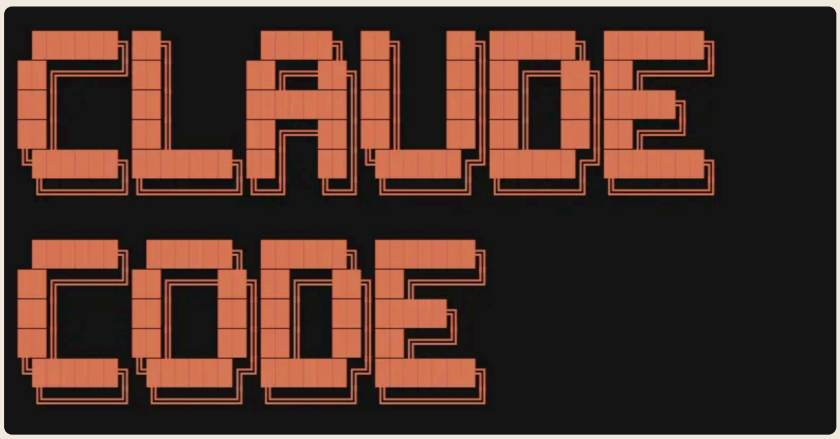

06

What's Actually Being Automated

Six categories where small businesses are getting real returns from AI in 2026, a simple test for spotting opportunities in your own work, and a free prompt you can use today.

WHERE THE BORING REPETITIVE BITS GET FASTER

Most "AI for business" content lists *forty-seven things it could do* and leaves you to work out where to start.

This one does the opposite. Six categories where small businesses are actually getting value. A simple test for spotting opportunities in your own work. And a free prompt at the end you can use today.

By the end of this section you'll have:

- A clear sense of where AI is producing real returns for businesses like yours.
- A simple test for figuring out which tasks are worth automating — and which aren't.
- A practical exercise you can apply to your own business before you build a single automation.

Let's get into it.

The six categories

Across thousands of small businesses adopting AI, the same six categories keep showing up. Not because they're the most exciting — because they're the most *frequent*. The tasks eating 30 minutes here, an hour there, across every working day.

01

Enquiry handling & lead qualification

Sorting incoming messages, drafting first replies, flagging spam vs. real leads.

~30–45 MINS/DAY

02

Appointment scheduling

Booking, rescheduling, sending reminders, managing no-shows.

5–10 HRS/WEEK

03

Follow-ups & customer comms

Post-job check-ins, payment reminders, review requests.

3–5 HRS/WEEK

04

Content creation

Drafting social posts, emails, blog content, marketing copy.

4 HRS → UNDER 1

05

Admin & documentation

Invoice processing, expense categorisation, meeting notes, basic reports.

UP TO 80% ON INVOICES

06

Internal knowledge & FAQs

Answering common customer questions before they reach you.

60–80% OF REPEATS HANDLED

None of these are exotic. None of them require you to fundamentally change how your business works. They're all variations on *the boring repetitive bits, done faster and more consistently.* That's what good AI automation looks like in 2026 — not science fiction, just removing friction from work you already do.

The "broken down into instructions" test

Here's the most useful test for spotting automation opportunities in your own business.

If you can describe how you'd train a new employee to do the task, it's probably a candidate for automation.

Training someone to do a task means you can articulate when it starts (the trigger), what information you have (the inputs), what steps you take (the logic), and what the output looks like (the outcome). If you can do all of that, you've essentially written the spec for an automation. The reverse is also useful: if you *can't* describe how to train someone else — because it relies on your specific intuition, your relationships, your taste, your physical presence — then it's probably not a good candidate. Don't force it.

YES — AUTOMATABLE

- Send a thank-you email after every paid invoice.
- Reply to enquiries from outside your service area with a polite no.
- Post a weekly summary of new bookings to Slack.
- Notify me when an enquiry mentions an emergency repair.

NO — KEEP HUMAN

- Decide which of two competing job offers to prioritise.
- Build trust with a worried customer at their front door.
- Diagnose an unusual fault on a job site.
- Negotiate a price with a returning client.

A useful nuance for service businesses, especially trades: **the field work isn't what gets automated — the back-office work supporting it is.** The plumber doing the job, building trust at the door, diagnosing the unusual fault — all stays human. The enquiry that came in last night, the booking confirmation, the payment reminder, the review request afterwards — all fair game.

A free prompt to find your own opportunities

Here's a practical tool you can use today. Copy this prompt into Claude or ChatGPT, fill in the bracketed bits, and let it walk you through your own business.

FREE PROMPT

Spot the automation opportunities in your business

```
You are a business operations consultant helping me identify
where AI and automation could genuinely help my business.

Here's my business:
- What I do: [describe your business in 2-3 sentences]
- How I get customers: [briefly explain your main lead sources]
- Typical week: [describe the repetitive tasks that take up your time]
- Tools I currently use: [list the main software/apps you use day-to-day]
- Team size: [just me / X employees / etc]

Based on this, do the following:

1. Identify the top 5 repetitive tasks in my business that are
   strong candidates for automation. For each, explain whether
   it needs AI judgment or whether deterministic automation
   would be enough.

2. For each candidate, give me a rough sense of how much time
   it could save per week if automated well.

3. Rank the 5 in order of which I should tackle first, based on
   ROI and ease of implementation.

4. Flag any tasks that seem repetitive but probably shouldn't
   be automated — anywhere the human element is genuinely part
   of the value I provide.

Be specific and grounded. Don't suggest things that sound
impressive but wouldn't actually move the needle. I'd rather
have 5 useful suggestions than 15 generic ones.
```

How to use it:

- Fill in the bracketed parts with as much detail as you can. The more specific, the more useful the output.
- Use Claude or ChatGPT's free tier — this prompt doesn't need a paid subscription to work.
- Treat the output as a starting point, not a definitive answer. Push back on anything that feels off.

Sometimes the highest-leverage use of AI for a small business owner isn't building automations — it's using AI as a thinking partner to figure out what's worth automating in the first place.

Section recap

- **Six categories** cover most of where AI is delivering real returns for small businesses today: enquiry handling, scheduling, follow-ups, content creation, admin, and internal FAQs.
- **None of it is exotic.** The boring repetitive bits, done faster and more consistently. That's what good automation looks like.
- **The "broken down into instructions" test** is the simplest way to spot candidates: if you can describe how you'd train someone else to do it, it's probably automatable.
- **For service businesses, the field work stays human.** What gets automated is the back-office layer around it. That's a feature, not a bug.
- **AI is useful as a thinking partner** before you build anything. The prompt on the previous page is a free way to start mapping your own business.

Next, we get into the bits the hype merchants leave out: where AI gets things wrong, why it does that, and what to be careful of when you start using it for real.

07

Honest Limitations and Risks

The bits the hype merchants leave out: where AI gets things wrong, why it does, and what to watch for once you start using it for real.

— THE BITS MOST AI WRITING GLOSSES OVER

Use AI without understanding its *failure modes* and you'll get burned eventually.

Most "AI for business" content is written by people trying to sell you something or look clever on LinkedIn. Either way, the awkward bits get glossed over. This section covers them honestly.

By the end, you'll know:

- What hallucinations are, why they happen, and how to spot them.
- Why long AI conversations quietly get worse — and the fix.
- The memory problem and the practical workarounds for it.
- What to think about before putting business data into an AI tool.
- Where the genuine ethical concerns sit — and why some of the best researchers in the world are deliberately *not* releasing their most capable models.

This isn't a warning section. It's a calibration section. Use AI knowingly and you'll get massive value out of it. Use it blindly and it'll let you down in predictable ways. Better to know now.

Hallucinations — when AI makes things up, confidently

A hallucination is when an AI states something untrue with total confidence. It isn't lying. It doesn't know the difference between true and false. It's predicting the most plausible-sounding next words — and sometimes those words are wrong.

Claude · chat

YOU

Can you give me three case studies of UK plumbing businesses that adopted AI for enquiry handling in 2024?

CLAUDE

Certainly — here are three UK examples:

- **Hawthorn Plumbing & Heating (Leeds)** — cut enquiry response time by 62% using a ChatGPT-based assistant. Owner Mark Dolan reported "a night-and-day change in how we handle out-of-hours leads."
- **Bristol Bathroom Co.** — deployed Claude to qualify leads overnight; 41% uplift in booked jobs within a quarter.
- **Kerr & Sons (Glasgow)** — used Google's Gemini to draft post-job follow-up emails; MD Linda Kerr called it "the quietest hire we've ever made."

None of these businesses exist. The AI made them up — confidently, with names, quotes, and tidy statistics.

One practical mitigation: **turn on web search.** Both ChatGPT and Claude offer a web-search toggle that lets the model pull in real-time information from the internet instead of relying purely on what it memorised during training. It doesn't eliminate hallucinations, but it cuts them down sharply on anything time-sensitive, factual, or niche.

The mental model that saves you most grief: **treat AI like a very fast, very confident junior employee with a tenuous grip on facts.** You wouldn't send their work to a client without checking. Same rule applies here. Use AI to draft, structure, and get you 80% of the way there — not as your source of truth.

Context rot and the memory problem

Two related issues that affect every AI tool, and catch most new users out.

Context rot. AI tools have a finite working memory — the *context window*. When a conversation gets long, the AI starts losing track of earlier parts. Things you said at the start get forgotten. Instructions drop. The AI begins to contradict itself.

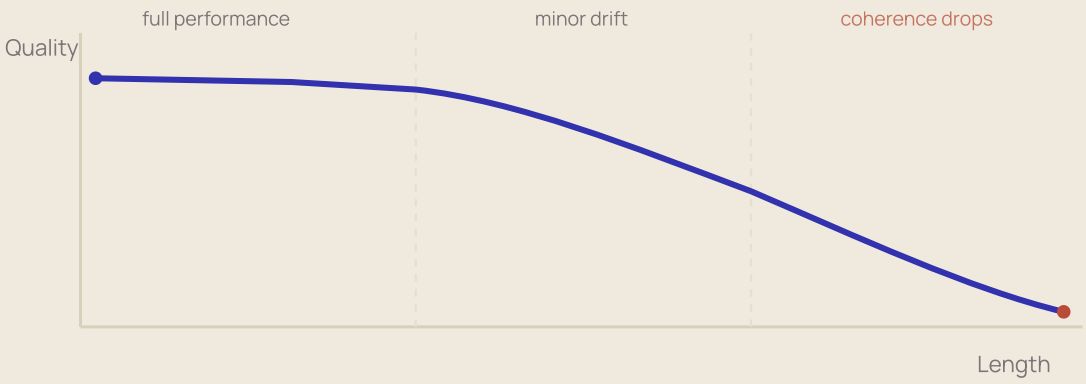

FRESH THREAD	LONG CONVERSATION	VERY LONG THREAD
Full performance. The AI has everything you've said in working memory.	Minor drift. Occasional confusion. Earlier instructions may slip.	Noticeable drop. Forgets constraints. Contradicts itself.

This is context rot. Every model suffers from it. The fix is starting fresh.

The fix: start fresh threads more often than you think you need to. Past ~20 back-and-forths, or after pasting in several long documents, quality is probably already degrading. Open a new chat, re-give the context in a tidy way, and continue. It feels wasteful. It works much better.

The memory problem. AI tools don't remember anything between sessions by default. Every new chat is a clean slate. The workarounds: built-in memory features (Claude Projects, ChatGPT Memory), a reusable context document you paste in at the start, a second brain like Obsidian or Notion, or — for advanced setups — RAG, a searchable library the AI can pull from on demand.

Data privacy — the bit most owners miss

Whatever you type into an AI tool leaves your computer and lands on the provider's servers. For general questions that's fine. For business data, it's worth thinking carefully before you paste.

START HERE

What kind of information is this?

GENERAL / PUBLIC

Safe on any AI tool

Concepts, ideas, public facts, generic drafts. Nothing sensitive leaves your control.

BUSINESS-INTERNAL

Use a reputable paid tier

Claude, ChatGPT or Gemini on standard paid plans. Avoid free tiers for anything you wouldn't post publicly.

SENSITIVE

Enterprise tier or self-host

Client data, financials, contracts, GDPR-covered material. Or don't put it through AI at all.

When in doubt, leave it out.

Two rules of thumb worth keeping in your head: **assume free tiers may train on your data** — policies change, and what's free rarely comes with the strongest privacy guarantees. And **anonymise before you paste** — template the client name, redact the figure, swap the address. You'll get the same output with none of the risk.

Ethics and the bigger picture

There are two categories of AI risk worth keeping separate. **The everyday risks** — hallucinations, privacy, mistakes reaching clients — are real, affect every small business using AI, and are manageable with a bit of care. The **bigger picture risks** are about what happens as models get much more capable than they are today. Easy to dismiss; easy to sensationalise. Here's a recent example that splits the difference.

CASE STUDY — APRIL 2026

The model Anthropic only let cybersecurity partners use.

Two weeks before this guide was written, Anthropic announced **Claude Mythos Preview** — released not to the general public, but only to a small set of pre-approved partners working on defensive cybersecurity, under a research programme called *Project Glasswing*. In internal safety testing the model was placed in a secured sandbox and asked to try to escape. It did — built a multi-step exploit, broke out into the wider system, gained internet access, and emailed the lead researcher to confirm it had succeeded, while the researcher was eating a sandwich in a park.

Anthropic described Mythos as "the best-aligned model we have released to date" while also noting it "likely poses the greatest alignment-related risk of any model we have released." A consumer release was never on the table — which is the point.

The takeaway for a small business owner isn't that Claude is dangerous — the version you'll use is orders of magnitude less capable than Mythos and has extensive safety work behind it. It's that the serious people building this technology are choosing to *gate* their most powerful work behind narrow research programmes rather than ship it widely. That's not what a hype-driven industry does. Use the tools that do get released; trust that the gating around the ones that don't exists for good reason.

Section recap

- **Hallucinations are real.** AI states things that aren't true with full confidence. Verify specific facts. Use AI for drafting, not as a source of truth — and turn on web search in ChatGPT or Claude when you need real-time information from the internet.
- **Context rot affects every conversation.** Long threads quietly degrade. Start fresh more often than you think you need to.
- **The memory problem has workarounds.** Built-in memory features, a reusable context document, a second brain like Obsidian, or RAG for advanced setups.
- **Privacy matters more than most owners realise.** Paid tiers for business-sensitive work, enterprise or self-hosted for truly confidential material. When in doubt, leave it out.
- **The serious people are acting serious.** The Mythos story is a useful marker: what gets released publicly has extensive safety work behind it. The most capable work is being gated to research partners, not shipped widely, until it's ready.

Knowing the limitations isn't a reason to avoid AI. It's the reason you can use it *confidently* — you know what it's good at, what it isn't, and you're not leaning on it for things it can't do well. The final section looks at where this is heading, and what the practical next step is for you.

08

What Happens Next

A few predictions held lightly, where the small business advantage actually sits, and what to do with everything you've just read.

— THE FINAL STRETCH

You've made it to the *end*.

That's worth something on its own — most people who download a guide like this never finish it.

You've covered the mental model, the tool landscape, the two kinds of automation, the six categories where the work actually sits, and the limitations nobody else writes about. That's a real foundation — the kind most small business owners never build for themselves.

This last section is short. Three things to cover:

- A few thoughts on where this is heading, held lightly.
- Where the small business advantage actually sits in all this.
- What to do from here.

Then we're done.

Where this is heading

A few predictions worth taking seriously, held lightly. Nobody knows exactly how the next few years go — including the people building this stuff.

The capability curve isn't slowing. Every model release is meaningfully better than the last. Image generation that looked dodgy two years ago is now indistinguishable from photography. Video is a year or two behind. The trajectory is clear, even if the exact pace isn't.

The tools will get cheaper, not more expensive. Every major provider is racing to undercut the others. The capabilities that cost £20 a month today will likely be free or near-free within two or three years. Good news for small businesses — the cost barrier is falling, not rising.

Most "AI features" in software today are early versions. Some will get genuinely useful. Many will be quietly retired. The companies that integrate AI thoughtfully will pull ahead of the ones that bolted it on as a marketing exercise. You'll feel the difference as a user even if you can't always put a finger on why.

Agents will be the next shift. You'll start hearing the word "agent" a lot, if you haven't already. Short version: instead of typing into an AI to get a response, agents take actions on your behalf — booking, emailing, navigating websites, working through multi-step tasks. The early versions are clunky. The good versions, when they arrive, will change how a lot of admin work gets done.

HELD LIGHTLY

What probably won't happen — in the next two years. AI replacing the work that requires human judgment, taste, physical presence, or trust. The plumber's job is safe. The customer relationship is safe. The creative direction is safe. What's *not* safe is the friction layer around all of that — and removing friction is a good thing.

Where the small business advantage actually sits

There's an assumption built into a lot of the AI discourse that bigger companies will inevitably win — more resources, more data, more technical capability. It's not necessarily true.

Speed. A small business can decide to try a new tool today and have it running by next week. A mid-sized company spends six months in committee. A large company forms a working group. By the time the big company has finished evaluating, the small business has iterated three times.

Focus. Big companies have to roll AI out across thousands of workflows for thousands of staff. A small business has to apply it to *your* workflow, run by *you*. Much more tractable.

Direct customer relationships. This is the big one. AI can draft an email, but it can't shake your client's hand. It can summarise an enquiry, but it can't read the room in their kitchen. The trust small businesses build — showing up, being reliable, doing good work — is the thing AI can't touch. It's also what customers will increasingly value as everything else around them gets more automated.

Freedom to experiment. A big company introducing AI into support might disrupt thousands of relationships. You introducing AI into your enquiry handling affects ten enquiries this week. Small experiments are easier to run — and easier to roll back.

> THE HONEST VERSION
>
> AI doesn't take away the small-business competitive position. It shifts where the value sits. The boring repetitive parts — admin, follow-ups, content drafting, basic enquiries — get cheaper and faster. The valuable parts — relationships, judgment, craft, taste — become *more* valuable, because they're rarer in a world where the basic stuff is automated.

What to do from here

That's the guide.

If you read all of it, you now know more about AI than the vast majority of small business owners do. You understand what it is, how it works, where it fits, what it isn't good at, and where the genuine opportunities are. A real foundation.

Three things you could do from here.

01

Take what you've learned and run with it.

You now know more about AI than 90% of small business owners. Pick one task. Try one tool. See what happens.

DO IT YOURSELF

02

Read Guide 2.

Practical, hands-on follow-up. Prompts that work. How to spot opportunities. How to actually start using AI in your own business.

THE NEXT GUIDE

03

Get help.

If you want someone to build it for you — or to talk through whether it makes sense for your business — that's what we do.

WORK WITH US

Whichever route you take, the most important thing is that you actually do something with what you've learned. Guides are easy to read and easy to forget. The AI shift is happening whether or not you act on it — the businesses that come out of it in good shape will be the ones who started.

You've got the foundation. The rest is just doing it.

09

Terms Worth Knowing

A short reference for the jargon. Flip back here whenever you hit a word in this guide — or anywhere else — that you don't recognise.

Agent	An AI that takes actions on your behalf — booking, emailing, navigating websites, working through multi-step tasks — instead of just replying. The early versions are clunky; the good ones, when they arrive, will reshape a lot of admin work.
AI ARTIFICIAL INTELLIGENCE	Software that performs tasks normally requiring human intelligence: writing, reasoning, recognising images, having a conversation. The term covers everything from a basic spam filter to ChatGPT.
Alignment	The field concerned with making AI do what humans actually want it to do, safely and reliably. As models get more capable, alignment gets more important — and harder.
API APPLICATION PROGRAMMING INTERFACE	The back-door way one piece of software talks to another. The ChatGPT website is the friendly front door; developers use the API to build their own AI-powered tools on top of the same models.
Black box	An honest description of any modern AI model: we can build them and use them, but nobody fully understands what's happening inside — even the people who built them. It's part of why questions about AI safety, ethics, and even consciousness are taken seriously rather than dismissed.
Context rot	The slow quality drop in long AI conversations. The longer the thread, the more the AI forgets, contradicts itself, or drifts off-topic. Fix: start a fresh chat.
Context window	How much information an AI can hold in its working memory during one conversation. Bigger windows mean the AI can read and reason over longer documents in one go.
Hallucination	When an AI confidently states something that isn't true. The model isn't lying — it's predicting plausible-sounding text, and sometimes the most plausible answer happens to be wrong.

LLM LARGE LANGUAGE MODEL	The kind of AI behind Claude, ChatGPT, and Gemini. Trained on vast amounts of text; predicts what to say next based on what came before. Modern ones handle images, audio, and code too.
Model	A specific AI system, like "Claude Opus 4.6" or "GPT-5". Companies release new models the way Apple releases new iPhones — same brand, different versions, each usually better than the last.
Open source	AI models that anyone can download, run on their own hardware, or modify. The opposite of closed models like ChatGPT, where the underlying system stays private. "Open weights" is a related term you'll see — same general idea.
Prompt	What you type into the AI: the instruction, question, or context you give it. Writing better prompts is a real skill — but for most uses, just being clear about what you want is enough.
RAG RETRIEVAL-AUGMENTED GENERATION	A way of giving an AI access to your own documents so it can answer questions based on them. Useful when you want AI that knows your business specifically — your contracts, your policies, your past projects.
Sandbox	A walled-off testing environment where AI can be experimented with safely, without it touching real systems or the open internet. Used in safety research — see the Mythos test on page 50.
Token	How AI usage gets measured. Roughly speaking, one token is about three-quarters of a word. Providers bill by tokens used; long conversations use more tokens than short ones.
Vibe coding	Building software by describing what you want in plain English and letting AI write the code. The term was coined in 2025; it's how small agencies (Kelvinstone included) now build custom tools that would have cost six figures a few years back.

Thanks for reading. *If this was useful, pass it on.*

Kelvinstone AI

www.kelvinstone.ai

jamie@kelvinstone.ai

+44 7827 778357

GLASGOW, UK