

aban news

# Anti-Hype

How solopreneurs put AI to work without the bullshit

*Expanded edition · 2026*

*Aban (Allen Chour) · [abannews.com](http://abannews.com)*

# Contents

- 01 Why this eBook exists
- 02 The 3-question bullshit filter
- 03 What AI is good at — and where it reliably fails
- 04 Your minimal AI stack
- 05 How to talk to AI so something useful comes out
- 06 Five routines that actually save time
- 07 Seven hype traps
- 08 Data protection in five minutes
- 09 Where to go from here
- 10 Templates to copy
- 11 A day, walked through
- 12 Tool by task
- 13 Plain-talk glossary: ten terms without fog
- 14 Sources & further reading

*Don't read it cover to cover like homework. Jump to the chapter that solves today's problem. The later chapters are the workshop: templates to copy, one fully worked day, and a tool-by-task overview.*

# Why this eBook exists

*Most AI guides sell you a feeling. This one doesn't.*

Since 2024 I've written a daily newsletter about AI for people who make money with it — not talk about it. Along the way I've tested hundreds of tools, deleted dozens, and noticed: most of what gets sold as a “2026 must-have” solves a problem you don't actually have.

This eBook is the distilled version for one person: the solopreneur whose day has 24 hours and who won't waste any of them on tool demos.

An example I see often: someone buys a 200-dollar subscription to an “AI content system” because a video promised 30 posts a week. Three weeks later there are zero posts — because the problem was never the writing, it was the missing decision about what to write at all. No tool on earth makes that decision for you.

## What you will NOT find here

- Prompts that supposedly “automate your business” but really just produce a nice paragraph.
- Tool lists with 50 entries, 47 of which you'll never open.
- The promise that AI makes you rich while you sleep.

## How to use this eBook

Short, self-contained chapters, each read in under ten minutes. Read the one that hits today's problem, try a single thing, and come back when the next question shows up. Knowledge you don't apply is just consumption that got more expensive.

**What you will find: a way of thinking that separates substance from show — and a few concrete routines that save hours.**

# The 3-question bullshit filter

Before you take an AI tool, course or trend seriously, run it through three questions. If it survives all three, it's worth a second look.

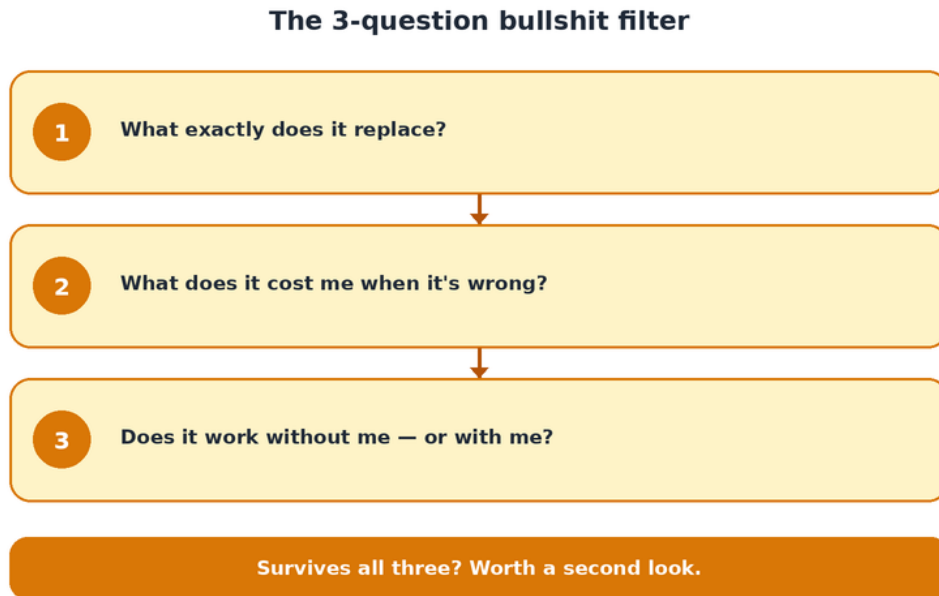


Fig. 1 — The 3-question filter at a glance.

## 1. What exactly does it replace?

Not “what can it do”, but: which concrete task you do manually today disappears? If the answer stays vague, it's show. If it's concrete, it's substance.

Test: describe the task in one sentence that starts with “Instead of X, I now do Y.” “Instead of typing every receipt by hand, I photograph it and have the data extracted” is concrete. “It makes my bookkeeping smarter” is marketing.

## 2. What does it cost me when it's wrong?

AI hallucinates. Always. The only question is how expensive a mistake is. Proofreading a draft post costs 30 seconds — delegate it. Copying a tax figure unchecked costs real money — verify every line.

Sort your tasks roughly into two buckets: reversible and expensive-irreversible. A headline that falls flat you fix in minutes — let AI hand you ten variants. A newsletter sent to 4,000 people with a wrong number can't be pulled back. The more expensive the mistake, the more of your own checking it deserves.

## 3. Does it work without me — or with me?

“Fully automatic” is almost always a lie. The tools that actually carry their weight are accelerators: you stay the head, AI does the dull work in between.

## **The most common thinking error**

Most people ask “what can this tool do?” and get impressed by the long answer. The better question is “which single task from my week does it reliably take off my hands?” One tool that truly does one thing beats ten that do everything a little.

**AI is an intern with a perfect memory and zero judgement. Treat it that way — and you'll rarely be disappointed.**

# What AI is good at — and where it reliably fails

*Before you put AI to work, you should know what it was built for — and what it wasn't. That spares you most of the disappointments.*

A language model is at its core a machine that predicts the next word. That sounds unspectacular, but it explains almost everything: why it phrases things brilliantly and invents facts in the same breath, why it never notices it's wrong, and why the same job gives slightly different answers twice. It doesn't compute truth, it computes probability.

## Where AI is strong today

- **Reshaping:** long text into short, bullet points into prose, formal into casual. Its signature discipline.
- **First drafts:** filling the blank page so you have something to edit instead of invent.
- **Sorting and summarising:** boiling large amounts of text down to the essentials.
- **Explaining language:** a technical term, an error message, a contract clause — in plain words.
- **Delivering variants:** ten headlines, five subject lines, three tones, and you pick.

## Where it reliably fails

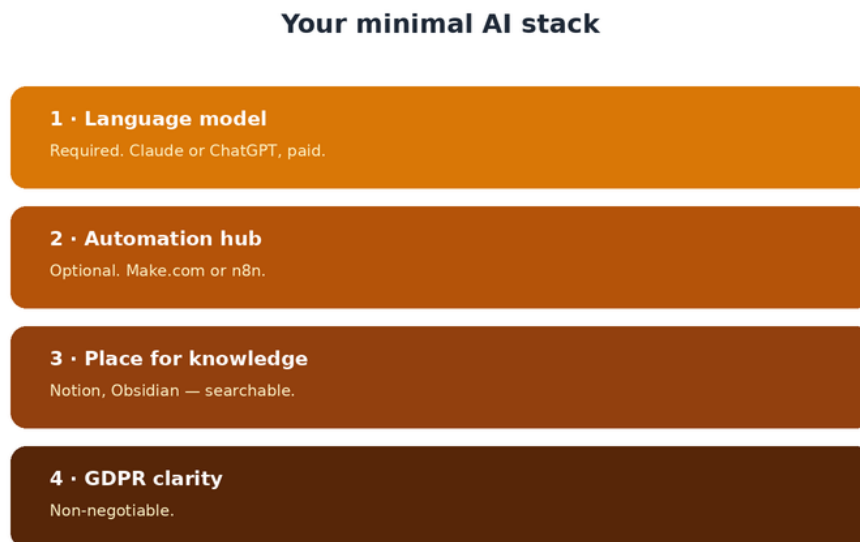
- **Facts and figures:** it cites sources, dates and statistics with full confidence, even when it just made them up.
- **Anything current:** the model only knows the world up to a cut-off date; what happened last week it doesn't know on its own.
- **Maths and counting:** on multi-step maths or counting it's wrong more often than you'd think.
- **Judging:** what is right, fair or strategically smart it does not decide — it only imitates how people write about it.

So the line doesn't run between easy and hard, but between language and truth. As long as it's about form, AI is a strong partner. The moment it's about facts, it's a claim machine you have to check.

**Use AI for anything you can proofread and correct yourself. Never use it as the final instance for something that has to be right.**

# Your minimal AI stack

*You don't need 20 subscriptions. You need four building blocks.*



*Fig. 2 — Four building blocks, no more.*

## 1. One good language model (required)

Claude or ChatGPT, paid tier. Your workhorse for writing, sorting, summarising. One is enough — learn it properly instead of constantly switching.

The roughly 20 to 25 euros a month for the paid tier is rarely the issue: if it saves you one hour a week, it has paid for itself on day one at any realistic hourly rate. What's actually scarce is your time to learn to drive it well — not the budget.

## 2. An automation hub (optional)

Make.com or n8n. This is where you glue services together. Only add it once you genuinely do a task manually three times a week.

## 3. A place for your knowledge

Notion, Obsidian, whatever — as long as it's searchable. AI is only as good as the context you give it.

Concretely: keep a simple text file with your audience, your tone, your three most common tasks and a few phrasings that sound like you. Paste that block at the top of every longer job. The difference between a generic answer and one that sounds like you is almost always just this context — not the model.

## 4. GDPR clarity (non-negotiable)

Before you dump customer data into a tool: where are the servers? Is there a data-processing agreement? A GDPR breach costs more than any productivity gain.

Do: anonymise what you don't need — “client A in industry B” is plenty for a draft proposal. Don't: paste whole mail threads with real names, addresses and contract details into some random free tool just because it's handy right now.

**Rule of thumb: process first, tool second. Whoever starts with the tool goes looking for a problem to fit it.**

# How to talk to AI so something useful comes out

*Most bad AI answers aren't the model's fault, they're the brief's. Four moves change almost everything.*

A good language model is like a capable intern: the clearer the brief, the less you fix afterwards. A vague brief gives you the average of half the internet — correct, but interchangeable. These four moves lift quality more than any tool switch.

## 1. Say who you are and who the text is for

“Write a text about taxes” is a dice roll. “I'm a tax advisor writing for self-employed tradespeople who are short on time — explain the small-business rule in plain words” gives you something usable. Context isn't decoration, it's half the brief.

## 2. Specify the format

Say what should come out: five bullet points, a table, a paragraph under a hundred words, an email with a subject line. Without a format you get prose you then have to take apart again.

## 3. Show an example

A single example of your style works harder than three paragraphs of explanation. Attach an old post, an earlier email or a phrasing that sounds like you, and say: “in this tone”. The model imitates — so give it something good to imitate.

## 4. Iterate instead of re-rolling

The first answer is a draft, not a finished product. Say concretely what's off: “too salesy”, “shorter”, “the second point is wrong, drop it”. Two or three targeted corrections get you further than ten rounds of “do it again”.

## The difference in one example

Weak: “Write me a newsletter about AI.” Strong: “Write the intro for my newsletter to German-speaking solopreneurs. Sober tone, informal, no hype, at most eighty words, the first sentence should spark curiosity without overselling. Topic: why most AI tools solve a problem you don't have.” The same effort to type — a night-and-day difference in the result.

**Context, format, example, correction. Whoever supplies these four needs no €47 prompt pack.**

# Five routines that actually save time

No magic — just routine work you delegate.

- **Pre-sort your inbox:** AI drops mails into three buckets — now, later, ignore. You still decide.
- **Turn one long-form into three formats:** one input, several outputs. You edit instead of starting from zero.
- **Write proposals faster:** a template plus the key facts, AI builds the draft. 45 minutes become 10.
- **Bundle research:** one clear brief with a sources requirement — then check the sources yourself.
- **Mirror your own week:** AI as a sober mirror, not a coach throwing confetti.

## One routine walked through

Take the proposal. Instead of starting with a blank document, you keep a fixed template: greeting, deliverable, scope, price, timeline, next step. You hand AI only the five key facts of this case plus the template and let it fill the draft. Then you read it once, fix the price and the tone — done. The thinking (what does this cost, what am I promising) stays with you; only the typing and formatting fall away.

Realistically that's closer to saving about half the time than the “10x” jump from the ads — and only after you've built the template cleanly once. It sounds unspectacular, but over a year it adds up to whole working weeks.

**Every routine keeps you as the final instance. AI delivers the raw material; your judgement makes the product.**

# Seven hype traps

- **Tool-collecting addiction.** Ten half-learned tools are slower than one you've mastered.
- **The autopilot myth.** “Runs by itself” usually means “breaks by itself when you're not watching”.
- **Prompt packs for €47.** A good prompt is a clear task — you can write it yourself, for free.
- **Vanity output.** Reach without substance is still emptiness.
- **AI as a strategy substitute.** A model won't decide for you who you want to be.
- **Blind trust in numbers.** AI confidently states statistics it invented.
- **FOMO as a business model.** Whoever scares you about missing out is selling you the cure too.

## How to spot a hype trap

Three warning signs almost always show up together: a deadline (“today only”), a round number with no source (“saves you 20 hours a week”) and a promise that takes your work away entirely instead of speeding it up. When all three sit in one ad, you can click it away fairly safely.

That doesn't mean ignore everything. It means wait for the second wave. What's still being shared after four weeks because it genuinely helps people — and not just because it was new — has earned your attention. The rest you can miss without harm.

**The uncomfortable truth: many who talk loudly about AI earn their money from the talking — not from the AI.**

# Data protection in five minutes

*GDPR isn't a reason to avoid AI — but a reason to think for a moment before the first click. Here's the check that takes five minutes and prevents expensive mistakes.*

The most common mistake isn't ill will, it's convenience: quickly dropping a client call into a free transcription tool, pasting a whole mail thread with real names because it's handy right now. This is exactly where the problems that get expensive later begin. Run these five questions once per tool — then you know where you stand.

## 1. Where are the servers?

Data processing in the EU is the easy path. With US providers you need a solid legal basis; many large tools now offer EU data residency or business tiers with better guarantees. If the page says nothing about it, that too is an answer.

## 2. Is there a data-processing agreement?

As soon as you process personal data, you need a data-processing agreement with the provider. Reputable providers offer it as a standard document. If you can't find one, no client data belongs in there.

## 3. Are my inputs used for training?

On many free tiers yes, on business tiers often not — but only if you actively set it or pick the right plan. Check that setting once, deliberately, instead of hoping.

## 4. What actually needs to go in?

The best data-protection measure is leaving things out. “Client A, construction sector, job X” is plenty for a draft proposal. Names, addresses and contract numbers you don't need for it — so leave them out.

## 5. Could I explain this to my client?

A simple test: if you could openly tell your client how you handle their matter with AI, it's usually fine. If you'd have to hide it, something is off.

**Anonymise what you can. A data-processing agreement where needed. And when in doubt, do the thing that really has to be right yourself. Here, data protection isn't a brake — it's protection for your business.**

# Where to go from here

If you've read this far, you already have the most important skill: patience for substance over the reflex for hype.

This eBook is a snapshot. AI changes weekly. What stays is the mindset: ask concretely, weigh the risk, judge for yourself.

## Your first step this week

Don't pick five things — pick one. Take the single task that felt most like dull repetition this week and try exactly one routine from chapter 4 on it. Roughly measure how long it took before and how long now. That one honest comparison gets you further than ten tool lists you've read.

That's exactly what I do every weekday in the newsletter — the three AI developments that affect your business today, in five minutes, no bullshit. No deadline, no pressure, unsubscribe any time.

*Subscribe for free: [abannews.com](https://abannews.com) — thanks for reading. — Aban*

# Templates to copy

*Seven building blocks you paste straight into your language model. Replace the placeholders in square brackets and it runs.*

One rule first: the more context sits at the top of the job, the less you fix afterwards. Always say who you are, who the text is for and in which tone — otherwise you get the average of the whole internet.

## 1. Summarise a long document

- **Summarise the following text in at most five bullet points.** Add one sentence on what I concretely need to do or decide. List open questions separately. Invent nothing that isn't in the text. Here is the text: [PASTE TEXT]

## 2. Turn notes into a post

- **Turn these raw notes into a LinkedIn post for [AUDIENCE].** Tone: sober, no hype, no chains of exclamation marks. The first line must spark curiosity without overselling. Max 1,200 characters. Notes: [NOTES]

## 3. Draft a proposal

- **Write the draft of a proposal.** Structure: greeting, the task in one sentence, deliverable, scope, price, timeline, next step. Factual, no sales talk. Key facts: client [INDUSTRY], task [TASK], price [AMOUNT], delivery by [DATE].

## 4. A cold email that doesn't read like spam

- **Write a short cold email to a [ROLE] at a company in [INDUSTRY].** Max 90 words, one concrete reference to their business, a single call to action (a 15-minute call). No flattery, no buzzwords. My offer: [ONE SENTENCE].

## 5. Shorten a text without losing the point

- **Cut the following text in half.** Keep the core message and the numbers, drop filler words and repetition. Same tone as the original. Text: [TEXT]

## 6. Find counter-arguments before others do

- **I claim: [YOUR THESIS].** Give me the three strongest counter-arguments and a short answer to each. Be hard, not polite — I want to see the gaps before a client finds them.

## 7. Pre-draft an email reply

- **Draft a polite, clear reply to the following email.** My goal: [WHAT YOU WANT TO ACHIEVE]. Short, friendly, without grovelling, same level of formality as the original. Email: [INCOMING EMAIL]

**Save the two or three templates you need weekly as text snippets. A template you can't find again is one you won't use.**

# A day, walked through

A realistic Tuesday for a freelance copywriter — once without, once with AI. No invented miracle number, just an honest sum.

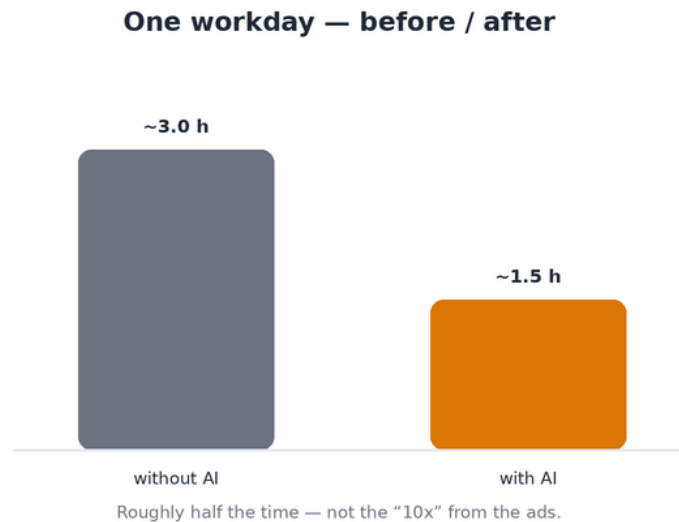


Fig. 3 — An honest halving, not “10x”.

Maria is a freelance copywriter with five clients. Her Tuesday has three recurring blocks: the inbox, one client proposal, one blog post from notes. Here is how it ran before, and how it runs now.

## Before: the manual Tuesday

- **Inbox (40 min):** read 30 mails one by one, sort, answer three.
- **Proposal (50 min):** blank document, rethink the structure, hunt for phrasings, format.
- **Blog post (90 min):** turn twelve bullet points into 800 words, reorder twice.
- **Total: about three hours**, much of it dull setup and formatting.

## After: the same Tuesday with AI

- **Inbox (20 min):** AI sorts into “now / later / ignore” and proposes three reply drafts. Maria reads, fixes, sends.
- **Proposal (20 min):** her fixed template plus five key facts, AI fills the draft. Maria checks price and tone.
- **Blog post (55 min):** AI builds a rough draft from the bullets, Maria writes the intro and ending herself and sharpens the claims.
- **Total: about an hour and a half** — the thinking stays with her, the setup falls away.

## What keeps the sum honest

That's roughly half the time — not the “10x” from the ads. And the jump didn't come on day one: Maria first needed a week to build her template and context block cleanly. She doesn't pour the saved time into more output, but into two things AI can't do: reaching out to new clients and editing texts more sharply.

Just as important: where a mistake would be expensive — the price in the proposal, the facts in the blog post — she still checks every line herself. The time saving comes from the setup, not from dropping the control.

**The gain isn't "more in the same time", it's "the same quality in less time — and the rest for what only you can do".**

# Tool by task

*Not “which tool is best”, but “what do I use for which task”. With honest notes on cost and GDPR.*

The names below are examples, not paid recommendations. Prices are rough orders of magnitude from early 2026 and change; check the current state before subscribing. There are no affiliate links here.

## Writing and editing

- **For:** text, summaries, rewriting, sorting.
- **Examples:** Claude, ChatGPT (both around 20-25 euros a month on the paid tier).
- **GDPR:** check EU data residency and a data-processing agreement; don't enter real client names without need.

## Research with sources

- **For:** an overview of a topic, with checkable links.
- **Examples:** Perplexity (free tier available, Pro around 20 euros a month).
- **Honest:** always open the cited sources yourself — “with sources” does not mean “correct”.

## Automation

- **For:** connecting services, recurring routines without click-work.
- **Examples:** Make.com (free quota, then from a few euros), n8n (open source, free when self-hosted).
- **Honest:** only worth it from three manual repetitions a week; before that it's tinkering.

## Images

- **For:** illustrations, simple graphics, social images.
- **Examples:** Canva (free tier, AI features on the paid plan from around 12 euros a month).
- **Honest:** check rights and trademarks; don't use generated images blindly as your own brand.

## Transcription and notes

- **For:** turning conversations, voice notes and calls into text.
- **Examples:** Whisper (open source, free locally), various web services (often from around 10 euros a month).
- **GDPR:** get consent for client calls; local processing is friendlier for data protection.

The rule above all: start with one good language model. Only when the same task there regularly takes too long is a second, specialised tool worth it.

**One tool you've mastered beats five you've merely installed. Extend the stack only when a concrete task forces you to.**

# Plain-talk glossary: ten terms without fog

*So no salesperson can intimidate you with jargon. Ten terms that come up in every other AI conversation — explained in one sentence.*

- **Language model (LLM):** the program behind ChatGPT or Claude that builds the most likely next word from text.
- **Prompt:** your brief to the model. The clearer it is, the better the answer.
- **Token:** the little chunks the model splits text into — roughly a word part. Billing and length limits run on these.
- **Context window:** how much text the model holds in view at once. When it's full, the beginning falls off the back.
- **Hallucination:** an invented but confidently delivered answer. Not a glitch, but the flip side of how it works.
- **Fine-tuning:** retraining a model on your own data. Unnecessary for most solopreneurs — good context in the prompt is enough.
- **RAG:** handing the model your own relevant documents before it answers, so it quotes from them instead of guessing.
- **Agent:** a model that doesn't just answer but carries out steps itself. Sounds big, breaks fast in practice — keep it narrow.
- **API:** the interface other programs use to talk to the model. Relevant only once you automate something.
- **Open source / local:** models you can run yourself. More control and privacy, but more effort.

**You don't need to master any of these terms to use AI sensibly. But whoever knows them won't be sold something they don't need.**

# Sources & further reading

*This eBook is experience, not a research paper. Where law and facts are involved, check the source yourself — here are the key ones, all official and free.*

## Law & data protection (for the data-protection chapter)

- **GDPR full text** — EUR-Lex, Regulation (EU) 2016/679. Search eur-lex.europa.eu for “2016/679”. Relevant for Art. 6 (legal bases) and Art. 28 (processors).
- **EU AI Act** — Regulation (EU) 2024/1689, also on EUR-Lex. Governs risk classes and transparency duties for AI systems.
- **Your data-protection authority** — the national/state DPA where you operate. First port of call for concrete questions.

## How the tools really handle data

- **Read the provider's own docs** — “Trust Center”, “Data Processing Addendum (DPA)” or “Privacy” on the provider's site. That's what counts, not a blog post about it.
- **Check the training setting** — in ChatGPT/Claude account settings, whether your inputs are used for training. Set it once, deliberately.

The rule of thumb stays: whatever has to be correct, you verify at the primary source — not with the AI and not in this book.