AI IS NOT A TOASTER

The stuff I wish somebody had told me when I started using AI.

No tech knowledge. No clever prompts. No jargon.

Total read time: about 15 minutes.

A few simple changes to how you use AI can make a massive difference to what you get from it. That's what this guide is about.

Mark Shaw
LinkedIn: Mark Shaw

AI is not a toaster

A toaster is simple.

Put the bread in. Push the lever down. Toast comes out. Job done.

AI can be simple too. In fact, once you get the hang of it, it's surprisingly simple.

But it doesn't work like a toaster.

When I first started using AI, that's pretty much how I treated it.

Question in. Push the button. Answer comes out. Job done.

And that's where I was going wrong.

The best results don't usually come from asking one question and accepting whatever pops out.

They come from talking to it. Questioning it. Correcting it. Asking it to explain itself. Telling it when it's misunderstood you. Getting it to ask you questions when you don't know what to ask.

None of that is difficult.

That's what I mean when I say AI is not a toaster.

You just need to stop treating AI like a machine that gives you an answer and start treating it like something you can have a conversation with.

So, if AI isn't a toaster, what is it?

I think of it less like a tool I'm operating and more like someone sitting across the table from me.

Someone I can bounce things off. Ask questions. Show things to. Challenge. Get a second opinion from. Tell them they've misunderstood me. Or just think out loud with.

A sparring partner, if you like. I'm still the one making the decisions.

There's just one slight problem.

The person sitting opposite me occasionally talks complete bollocks with enormous confidence. We'll come back to that. 😂

So, everything in this guide really comes down to three simple habits:

Talk to it. Challenge it. Check it.

That's the mindset. Everything else is just examples.

Now what?

I remember opening AI for the first time and staring at the empty box.

Now what?

What the hell am I supposed to type?

I wasn't technical. I'm still not. I don't write code. I've never owned a gaming console and I still occasionally print things out and file them.

So, I did what seemed obvious. I asked AI questions. It gave me answers.

Some were good. Some were incredibly generic. Some went on for three pages when three sentences would have done. And some, I later discovered, were wrong.

I'd made one fairly enormous assumption:

It's AI. It must know.

Nope.

It took me about a year of using AI, getting things wrong, arguing with it and generally messing about before the penny dropped. **Hopefully this guide saves you the year.**

It’s not about which AI you use

This isn’t a guide to ChatGPT. Or Gemini. Or Claude. Or whichever AI everyone is talking about by the time you’re reading this.

There are no screenshots showing you which button to press.

I’m not going to explain every feature. And I’m not going to tell you that you need a particular AI subscription.

This is about you, not the software.

Buttons change. Features change. AI companies change.

The basic habits in this guide don’t.

Stop Googling it

Google has trained us to reduce a problem to a few keywords.

Want to know something? Type a few words. Search. Results. Done.

So, when AI arrived, lots of us naturally did the same thing.

Question in. Answer out. Done.

Except that’s throwing away one of the best things about AI.

You can talk to it.

Explain what you’re actually trying to do. Read the answer. Don’t like it? Say so.

> “No. That’s not what I meant.”
>
> “Explain that in plain English.”
>
> “I don’t need all this. Just answer the question.”
>
> “I don’t think that’s right. Check it.”

You haven’t broken it. It won’t get offended. It won’t spend the rest of the afternoon sulking because you were rude about paragraph three.

And when I say talk to it, I mean it literally.

Most AI services let you use your voice. If staring at an empty box puts you off, don’t type. Just start talking. Explain the problem exactly as you would to another person.

Rambling a bit? Fine. Don’t know the right words? Fine. Change your mind halfway through? Fine.

Type or talk. Whichever feels more natural. The important bit is that you stop treating it like a search box.

Stop asking. Start having a conversation.

But I don't know what to ask

Good. Neither did I.

This might be one of the most useful things in this entire guide:

You don't need to know the right question.

Let's say you need a new TV. You know nothing about TVs. You don't know your OLED from your QLED and you don't particularly want to spend Sunday afternoon finding out.

Tell AI:

> "I need a new TV. I don't know enough about TVs to know what I should be asking. Ask me a few questions, one at a time, and help me work out what I actually need."

Now AI starts asking you.

What's your budget? How big is the room? What do you watch? Do you play games?

You answer normally. It asks something else. Suddenly you're having a useful conversation about what you actually need.

Bad spelling? Doesn't matter. Don't know the terminology? Doesn't matter. Can't explain yourself perfectly? Start badly.

You don't need to get better at talking to computers. The computer needs to get better at talking to you.

Make AI work your way

Here's a perfect first job for that trick.

Tell AI how you like to work.

You don't need to spend an afternoon fiddling with settings. You don't need three pages of personal information. And you definitely don't need to tell it your inside-leg measurement.

Let AI do the work. Tell it:

> "I'd like you to help me personalise the way you work with me. Don't give me a huge questionnaire. Ask me no more than five simple questions, one at a time."

Ask it to find out what you mainly want to use AI for, roughly how much you already know, how you like things explained, whether you want it to challenge you, and what you particularly like or dislike in AI answers.

Then ask it to create a short set of personalisation instructions you can use with your AI service.

Useful rules to include:

- Never invent facts, figures, names, quotes or sources.
- If you don't know, say so.
- If you're uncertain, tell me.
- Check important information when you reasonably can.
- Separate facts from assumptions and opinions.
- Don't agree with me just to be agreeable.

Some of those might seem odd right now. Why would you need to tell a computer not to make things up? Keep reading.

Will these instructions make AI perfect? No. They're instructions, not fairy dust. You still need to use your judgement.

Not sure where to put them? Your AI's own help pages will tell you. (Asking the AI itself is surprisingly hit and miss. More on that later.)

One more thing: lots of AI services now remember things about you as you go. Ask it:

> "What do you remember about me?"

You might be surprised. If something's wrong or out of date, tell it.

Show it. Don't describe it.

This one was a revelation for me.

You don't have to describe everything in words. Most AI services let you add a photo, a screenshot or a document.

The boiler is flashing a code you don't recognise. A letter you don't understand. A plant that's slowly dying on the windowsill. A form you're not sure how to fill in. An error message on your laptop.

Take a photo. Then ask:

> "What's this telling me, and what should I do?"

A picture saves you trying to explain something you don't have the words for. Which, if you're anything like me, is half the problem.

(Same common sense as everything else: think twice before photographing anything with bank details or somebody else's private information on it.)

Forget the magic prompts

If you've been anywhere near TikTok, YouTube or social media, you've probably seen them.

10 prompts you MUST know!

This SECRET prompt changes EVERYTHING!

You're using AI WRONG!

How I made £10,000 a week using AI.

Good for them.

This isn't that guide.

There's nothing wrong with giving AI good instructions. But you don't need to learn another language or memorise somebody else's magic words to use AI well.

And don't swing too far the other way. Once you learn a bit, it's tempting to overcomplicate everything.

Sometimes you're standing next to a rental car you've never driven before and you can't work out how to make the interior lights come on when you open the door.

You don't need a carefully constructed prompt. Ask:

> "I've got a Honda Fit rental car. How do I make the interior lights come on when I open the door?"

If AI needs to know the year, it can ask you. Or just send it a photo of the controls.

That's it.

Start simple. Add more detail when it's actually needed. Don't spend five minutes building the perfect prompt for something AI could have helped you with in thirty seconds.

Sometimes, just ask the bloody question.

And no, this guide won't teach you how to make £10,000 a week while lying on a beach.

If I ever work that one out, I'll let you know.

AI can make stuff up

I had absolutely no idea about this when I started.

I assumed:

Computer + Artificial Intelligence = knows the answer.

Not necessarily.

AI can give you a fact that isn't a fact. A quote nobody said. A statistic that doesn't exist. A book that was never written. A source that isn't real.

And here's what makes this particularly sneaky:

It doesn't necessarily look wrong.

It can be detailed. Professional. Confident. Beautifully written.

And complete bollocks.

There's a technical term for this: *hallucination*.

I prefer: **making stuff up.**

"Did you look that up?"

There's another catch. AI learns from a huge amount of information, but that learning stops at a certain date. Unless it searches the internet for you, it may not know about anything that's happened since. And it doesn't always make it obvious whether it searched or not.

So, for anything that changes (prices, opening times, rules, train times, who's in charge of what, anything "latest"), ask:

> "Did you look that up, or is that from memory?"

If it's from memory, ask it to search. Then check the source it gives you.

Test it on something you know

Want to see this for yourself?

Ask AI about something you know really well. Your job. Your town. Your hobby. Your football club's history.

You'll probably be impressed. And then, at some point, you'll spot something that's just… wrong. Said with total confidence.

That's the moment it clicks.

A wrong answer looks exactly the same as a right one.

Remember that feeling the next time it's telling you about something you know nothing about.

It worked yesterday

This one drove me mad for ages.

I'd ask AI to do something. Brilliant. I'd ask again. Brilliant. Ninety-nine times.

Then the hundredth time: rubbish. Same question. Same AI. Different result.

Why?

Because AI isn't like normal software. A calculator gives you the same answer every time. Two plus two is always four. AI doesn't work like that.

It writes a fresh answer every time. It isn't looking up a stored answer. It's writing a new one, with a bit of chance deliberately built in. Ask the same question twice and you'll get two different answers. Usually both fine. Occasionally not.

The AI behind the box keeps changing. AI companies update their systems all the time, often without telling you. Something that worked brilliantly last month may behave differently today.

You don't always get the same "brain". Some services quietly decide which version of their AI answers you. Sometimes a quicker one. Sometimes a more careful one.

Small things change the answer. Slightly different wording. What's earlier in the conversation. What it remembers about you. Whether it searched the web.

Think of it less like a calculator and more like asking a colleague the same question on Monday morning and Friday afternoon. Mostly the same answer. Not always.

Just because it worked yesterday doesn't mean it's right today.

Which brings us neatly to…

Treat AI like a six-year-old

Let me qualify that before someone writes me an angry email.

I'm **not** saying AI has the intelligence of a six-year-old.

I'm talking about **checking what it's done.**

Imagine you've got a six-year-old. Their bedroom looks like a small bomb has gone off.

You say: **"Go upstairs and tidy your room."**

Three minutes later they come charging back down.

“DONE!”

Really? That was quick.

You go upstairs. The room looks immaculate. For about five seconds.

Then you open the wardrobe. Everything’s been shoved inside.

You look under the bed. Clothes. Toys. Two books. A football. A plate. Something you don’t particularly want to identify.

Technically, the room looks tidy.

But the job you thought had been done hasn’t really been done.

AI can be exactly the same. Give it a long document and ask it to analyse it. A second later: **DONE!**

Bloody hell. That was quick.

And because the answer appears instantly and sounds intelligent, it’s very easy to assume it has done exactly what you asked.

So, when AI says “DONE!”

“I’ve checked it.”
Great. What did you check?

“I’ve analysed the document.”
Great. Show me where you found that. What page?

“This is definitely correct.”
Great. How do you know? What’s your source?

Asking AI to check itself is a good start. It isn’t proof. AI can confidently tell you it’s checked something and still be wrong.

If it matters, **check the actual source yourself**. If it quotes a document, find the quote in the document. If it gives you a source, open it and check it really says what AI claims.

You don’t need to interrogate AI every time you ask it for a birthday-card joke. Use some common sense.

But health? Money? Legal matters? Travel rules? Important documents? Something you’re about to repeat publicly as a fact?

Check the bloody thing.

The more it matters, the more you check.

AI sounding certain and AI being correct are not the same thing.

Get a second opinion

Here's a quick, free check for anything important.

Paste the same question into a different AI. ChatGPT, Gemini, Claude, Copilot. Whichever.

If they agree, that's reassuring. (Not proof. They can be wrong together.)

If they disagree, that's your cue to dig deeper.

Thirty seconds. Costs nothing. Catches a surprising amount.

Have a good argument with it

This completely changed how I use AI.

At first, I treated AI's answers with far too much respect.

Then I started arguing with it.

Much better.

Something sounds wrong? Say so.

Don't understand? Ask it to explain more simply.

It's misunderstood you? Tell it.

You disagree? Tell it why.

Sometimes AI will be wrong. Sometimes you'll discover you're wrong. Sometimes you'll both end up somewhere better.

The first answer isn't the answer. It's the start of the conversation.

But beware of the yes-man.

AI can be very agreeable.

Tell it: **"I've had a brilliant idea…"** and don't be too surprised when it finds lots of brilliant things to say about your brilliant idea.

Lovely.

But what if your idea is rubbish?

I don't want AI constantly telling me how clever I am. I've got enough friends for that.

Ask:

> "What's wrong with this?"
>
> "Why might this not work?"

> "What am I missing?"
>
> "What's the strongest argument against this?"

And if you push back and it suddenly changes its mind, don't assume the new answer is automatically better. Ask why it changed its mind.

Don't just use AI to prove yourself right.

Ask for options, not an answer

Instead of asking "What should I do?", try:

> "Give me three different ways to handle this. Tell me the pros and cons of each, then tell me which you'd pick and why."

Now you get a recommendation **and** you can see what you're choosing between.

And you're still the one making the decision. Sparring partner, remember.

Don't ask AI about its own buttons

This one surprised me.

AI is often wrong about itself.

Ask it where a setting is, what a button does or what it can do, and it can confidently send you to a menu that doesn't exist. Or one that moved three updates ago.

Which is quite funny when you think about it. And a perfect reminder that confident and correct are not the same thing.

Need help with the buttons? Check the company's own help pages, or ask the AI to search them for you and show you where it found the answer.

Start a fresh chat

Sometimes a conversation just goes wrong.

AI keeps repeating the same mistake. Gets muddled. Forgets something you told it earlier. Or just goes a bit weird.

Don't spend twenty minutes arguing with it. Arguing is good. Arguing in circles isn't.

Start a new conversation.

Long conversations get cluttered, and AI can lose track of things you said near the start. A fresh chat is a clean slate.

Handy trick: before you start again, ask:

> "Summarise what we've agreed so far in a short paragraph."

Paste that into the new chat. You keep the useful bits and lose the muddle.

Same goes for changing subject. New topic? New chat.

Your name is on it

If you send it, post it or hand it in, it's yours. Not AI's.

AI can draft a brilliant email, report or LinkedIn post. But don't just copy and paste it. Read it. Change it. Make it sound like you.

People are getting much better at spotting AI-written text. The over-polished phrases. The suspiciously perfect lists. The "I hope this email finds you well".

Want it to sound more like you from the start? Paste in something you've written yourself and say:

> "Write it like this. Same tone. Same sort of words."

AI does the heavy lifting. You have the final say.

A quick word about privacy

Use some common sense about what you share. Don't casually paste passwords, banking details, confidential information or other people's private information into an AI service. That goes for photos and screenshots too.

If something is sensitive, stop and think before you upload or type it.

And know when to stop

AI will nearly always find another improvement if you ask it. Another word to change. Another sentence to rewrite. Another version that's supposedly better.

Sometimes it genuinely is better. Eventually, you're just fiddling.

Learning when to say "that's good enough" might be just as important as learning how to ask AI to make something better.

That’s pretty much it

Seriously.

You don’t need a degree in AI. You don’t need a course. You don’t need to understand how the technology works. And you don’t need 500 secret prompts.

When you next see that empty box:

Start talking.

Don’t know where to start? Say so.

Don’t know what you should be asking? Get AI to ask you.

Easier to show it? Take a photo.

Don’t understand the answer? Tell it.

Too much waffle? Tell it.

Think it’s wrong? Challenge it.

Want another point of view? Ask it to disagree with you.

Gone a bit weird? Start a fresh chat.

Something important? Check it. Even if it worked yesterday.

Remember the toaster.

Don’t put the bread in, push the button and accept whatever pops out.

And when AI proudly announces: **“DONE!”**

Have a quick look under the bed.

STOP ASKING… START HAVING A CONVERSATION.

Your cheat sheet

Print this page. Stick it next to the computer. (I would.)

TALK TO IT

- “I don’t know enough about this to know what to ask. Ask me questions, one at a time.”
- “Here’s what I’m actually trying to do…”
- “Here’s a photo. What’s this telling me, and what should I do?”
- “Explain that in plain English.”
- “I don’t need all this. Just answer the question.”
- “Give me three options, the pros and cons, and which you’d pick.”
- “Write it like this.” (then paste in something you wrote)

CHALLENGE IT

- “No. That’s not what I meant.”
- “What’s wrong with this?”
- “What am I missing?”
- “What’s the strongest argument against this?”
- “Why did you change your mind?”

CHECK IT

- “Show me where you found that. What page?”
- “What’s your source?”
- “Did you look that up, or is that from memory?”
- “Prove to me you’ve actually done what I asked.”
- Ask a second AI the same question.
- If it matters, check the source yourself.

WHEN IT GOES A BIT WEIRD

- “Summarise what we’ve agreed so far.” Then start a fresh chat.

A bit about me

I'm Mark Shaw. I'm 60. I'm not a coder, an AI expert or particularly technical.

I've spent the last 18 months figuring out how AI can actually be useful in normal everyday life. Mostly by using it, getting things wrong, questioning it and occasionally arguing with it.

Before AI, I spent ten years explaining Twitter to businesses. I spoke across the UK and Europe, trained hundreds of businesses and did more than 230 TV and radio interviews, including the BBC and Sky News.

So, I'm not claiming to be an AI expert.

I'm quite deliberately coming at this as a normal bloke who started knowing bugger all about it.

This guide is simply the stuff I wish somebody had told me at the beginning.

Know someone who's still staring at that empty box?

Please feel free to send them this free guide.

No catch. No sign-up. No secret £10,000-a-week AI system waiting on the next page.

Just send it to a friend who's not sure how to start with AI.

It might save them the year it took me.

Mark Shaw
LinkedIn: Mark Shaw