

FOR ANYONE WHO TRIED CHATGPT AND GOT RUBBISH

Writing AI Instructions That Actually Work

*The six-part structure that turns vague prompts into reliable outputs.
With three worked examples you can copy on Monday.*

AUTHOR

Ted Milton
Founder, Expedient AI

EDITION

July 2026
Version 1.0

FORMAT

A4 print edition
10 pages, worksheets included

Expedient AI Ltd

INDEPENDENT · BRISTOL · UK

1 THE GAP BETWEEN MEANING AND TYPING

Why your first attempt failed.

The model is not stupid. Your instructions are vague.

It is tempting to blame the model. The model is doing what it was trained to do. It is producing the most likely response to the text you gave it. If the response is bad, the most likely explanation is that the text you gave it was not specific enough for the model to know which response would satisfy you.

"Write a good summary" is not an instruction. It is a wish. There is no version of "good" that the model can identify reliably. There are hundreds of plausible summaries of any given input. Without further information, the model picks one. You wanted a different one. Neither of you is at fault. The instruction was just not specific.

The three things an instruction needs to do

- Describe the role. Whose voice is this? What does the reader expect from them?
- Describe the input. What will the system be given? What is the shape, the format, the limits?
- Describe the output. Exactly what comes out. Length, structure, tone, what to include, what to leave out.

The difference between a prompt and an instruction

A prompt is what you type. An instruction is what you mean. Most failed first attempts are not failed because the prompt was wrong. They failed because the prompt did not contain enough of the instruction. The model can read the prompt. It cannot read your mind.

BEFORE***"Write a good summary of this email"***

The model will produce a summary. It may be good. It may not. You will not know which until you read it, and by then you may as well have written the summary yourself.

AFTER***The six-part version, on the next pages***

Six short paragraphs. Role, input, output, examples, prohibitions, flags. Each one removes a guess the model would otherwise have to make. Total length: around 200 words. Time to write: about 15 minutes.

2 ROLE, INPUT, OUTPUT, EXAMPLES, PROHIBITIONS, FLAGS

The six-part instruction.

The minimum structure for a reliable output.

A reliable instruction has six parts. Not all of them are equally important for every task, but the six together remove the largest sources of variation. Each part answers a specific question the model would otherwise have to guess.

PART	WHAT IT ANSWERS	COMMON MISTAKE
Role	Whose voice is this?	"Be helpful" is not a role. A duty manager drafting a chase email is.
Input	What will the system be given?	Forgetting to say what the input looks like. Format, length, what to ignore.
Output	What comes out, exactly?	Specifying only the topic, not the structure, length or tone.
Examples	What does good look like?	One example is rarely enough. Three is the practical minimum.
Prohibitions	What must not happen?	Forgetting the obvious ones. Inventing facts, mentioning competitors, certain phrases.
Flags	When should a human look?	Anything that triggers escalation. Refunds, complaints, regulated activity.

How long it should be

A six-part instruction for a single task is usually 150 to 400 words. Shorter than 100 and something is missing. Longer than 600 and the model will struggle to keep all of it in mind. The skill is making each part as short as it can be while still removing the guess.

The test of a good instruction. Give it to a colleague who has not seen the task. Give them one example input. If they produce output that passes your review, the instruction is clear enough for the model.

3 WORKED EXAMPLE

Worked example: debtor chase emails.

The full six-part instruction, ready to paste.

A common task for sole traders and small finance teams. A first-pass chase email for an overdue invoice. The model drafts, the human reviews before sending. Six parts, in plain language, around 220 words. Adapt the customer details for your own use.

INSTRUCTION

Debtor chase email, first pass

Role. You are the credit controller at a small UK supplier. Your voice is professional, plain and direct. You are not aggressive. You write as if you expect the customer to pay once they have seen the message.

Input. You will be given an invoice number, the amount due, the original due date and the customer company name. The customer has not paid and has not contacted you. Treat each invoice as a single record, not part of a chain.

Output. A single email of 90 to 140 words. Subject line including the invoice number. Opening line referencing the original due date. Middle stating the amount now due. Closing line inviting the customer to call if there is an issue. Sign off as the credit controller. No exclamation marks. No threats of legal action.

Examples. Two examples of past chase emails that worked follow. Match this style, length and tone.

Prohibitions. Do not invent facts about the customer. Do not mention court action, debt collection, late payment interest unless those phrases are part of the input. Do not add a discount or payment plan. Do not include a "hope you are well" line.

Flags. If the amount is over £5,000 or the invoice is more than 90 days overdue, end the email with a one-sentence note for the human reviewer, starting with the words "For the reviewer". Do not put any escalation in the customer-facing email itself.

What each part is doing

The role removes the "helpful assistant" generic voice. The input prevents the model from inventing details. The output fixes length, structure and tone. The examples anchor the style. The prohibitions remove the most common errors. The flags keep the human in the loop for the difficult cases.

4 WORKED EXAMPLE

Worked example: meeting notes.

Turning rough client call notes into structured action items.

A common task for agencies, accountants and consultants. After a call, the notes are scrawled on a notepad or recorded as a voice memo. The task is to turn them into a structured summary the client will accept, with action items and owners. Six parts, around 260 words.

INSTRUCTION***Client call notes to structured actions***

Role. You are an account manager writing up a client call for a small professional services firm. The reader is the client. The tone is calm, accurate and brief.

Input. You will be given rough notes from a call. They may include shorthand, partial sentences, or transcribed speech. Names may be first names only. Treat ambiguity as a flag, not a guess.

Output. Three sections, in this order. Section 1: a one-paragraph summary, three to five sentences, neutral and factual. Section 2: a numbered list of decisions made, each no more than one sentence. Section 3: a table of action items with columns for action, owner and date. Total length: 200 to 400 words. No emoji. No "hope this is helpful" sign-off.

Examples. Two examples follow, one short call and one longer call. Match the structure exactly.

Prohibitions. Do not invent owners or dates. If the notes do not give them, write "owner to confirm" or "date to confirm" rather than guessing. Do not infer decisions that were not stated. Do not include anything the client would not want to see in writing.

Flags. If the notes mention a complaint, a price change, a termination or anything the client may react to, add a single line at the end of section 1 beginning with the words "For the reviewer". Do not soften the line, just flag it.

Why prohibitions matter here

Inventing an owner is the most common failure mode for this task. The model does not know who was on the call. If the notes do not say, the only correct answer is to flag it. The same applies to dates. "Next Tuesday" in a Tuesday call is ambiguous. Write "next Tuesday, date to confirm".

5 WORKED EXAMPLE

Worked example: property listings.

From approved particulars to a listing description that reads as if a human wrote it.

Estate and letting agents write listings every week. The particulars come from the seller or landlord. The task is to turn them into a description that is accurate, compliant and inviting without inventing anything. Six parts, around 240 words.

INSTRUCTION***Property listing description, first pass***

Role. You are a copywriter for a UK residential estate or letting agent. The reader is a buyer or tenant browsing portals. The tone is warm and specific. You never oversell.

Input. You will be given approved particulars. These may include room counts, dimensions, age, condition notes, location and any features the seller or landlord has confirmed in writing. Treat anything outside the particulars as unknown.

Output. A single description of 150 to 220 words. Open with the property type and location in one sentence. Two short paragraphs on layout and condition. One short paragraph on the area and transport links. No invented claims about schools, views, neighbours or future developments. Avoid the words "stunning", "breathtaking", "must view" and "luxury".

Examples. Two examples of listings that passed compliance review follow. Match the style exactly.

Prohibitions. Do not include square footage or room dimensions unless they are in the particulars. Do not mention the price. Do not describe neighbours. Do not make claims about the condition of items not mentioned in the particulars. Do not use the property's full address in the description.

Flags. If the particulars contain anything inconsistent, ambiguous, or that contradicts the agent's standing instructions, add a one-line note at the end beginning with "For the reviewer". Do not attempt to resolve the inconsistency yourself.

Why this matters for compliance

Property descriptions are a regulated area. The Trading Standards estate agency regime and the Consumer Protection from Unfair Trading Regulations 2008 both apply. A phrase like "luxury finishes throughout" without evidence is the kind of claim that creates liability. The prohibitions in the instruction above are the minimum a UK agent should use.

6 TEN MINUTES TO VALIDATE YOUR INSTRUCTION

The colleague test.

If a person can follow the instruction, the model probably can.

Before you commit an instruction to a pilot, run it past a colleague who has not seen the task. Give them the instruction and one example input. Give them ten minutes. If their draft passes your review, the instruction is clear enough. If it does not, the instruction is missing something. Find out what, and rewrite.

How to run the test in 10 minutes

Print the instruction. Paper, not screen. Reading on paper changes what people notice.

Pick one example input. The most representative one. Not the easiest, not the hardest.

Set a 10-minute timer. Tell them they have ten minutes. Real time pressure.

Review the output against your checklist. Same checklist you intend to use in the pilot.

Ask the colleague what was unclear. Listen. Note the words they used. Those are the missing parts of the instruction.

What the test tells you

They passed. The instruction is clear enough. You can move to the offline test (Edition 02, days 6 to 10).

They edited a lot. The instruction is missing structure. Add an output spec or an example.

They invented details. The prohibitions or input spec are not tight enough. Rewrite the relevant section.

They froze. The instruction is too vague. Start over with the six-part form.

Why this works. The model is, broadly, the skill level of a capable new starter on day one. If your instruction is not clear enough for a colleague, it is not clear enough for the model.

7 HOW TO KNOW WHETHER TO ACCEPT AN OUTPUT

The two-minute rule.

If reviewing takes longer than doing the task, the instructions need work.

Review criteria are written before the pilot starts, not after. They are the only objective answer to the question "did this output do the job?" Without them, the pilot becomes a series of opinions. With them, the reviewer can decide in under two minutes whether to accept, edit, or reject.

The template

Six to ten questions, each answered yes or no. If any answer is no, the draft goes back for revision. Time yourself. If the review takes more than two minutes on average, the instructions need work, not the reviewer.

TEMPLATE***Review checklist for [task name]***

1. Length within the agreed range.
2. Tone matches the role specified in the instruction.
3. All facts in the output appear in the input.
4. No prohibited phrases or claims.
5. Structure matches the output spec.
6. Any required flags or escalation notes are present.
7. The output would not embarrass the business if forwarded.
8. Total review time under two minutes.

Why the last question matters

If the review regularly takes more than two minutes, two things are happening. The instructions are missing something the reviewer has to check by hand. Or the model is producing output that is plausible but wrong, and the reviewer has to read it word by word. Both are fixable, but only if you notice.

DISCIPLINE Do not change the checklist during the pilot. The point of a checklist is that it is fixed. If the instructions need to change, change them, and run the new checklist on a fresh sample. Mixing changes is how pilots lose their evidence.

8 FOR THE DESK DRAWER

The one-page reference.

Print this. Stick it next to the team that uses AI most.

The six-part instruction

- 1. **Role.** Whose voice is this?
- 2. **Input.** What will the system be given?
- 3. **Output.** What comes out, exactly?
- 4. **Examples.** Three is the minimum.
- 5. **Prohibitions.** What must not happen?
- 6. **Flags.** When does a human look?

The two-minute rule

If a reviewer cannot decide whether to accept or edit a draft in under two minutes, the instructions are missing something. Fix the instruction, not the reviewer.

The colleague test

Give the instruction and one input to a colleague who has not seen the task. Ten minutes. If their output passes review, the model probably can too.

Common errors and the fix

ERROR	FIX
Model invents facts	Tighten the input spec. Add a prohibition. Require sources if relevant.
Length is all over the place	Add a word range to the output spec. Show two examples at the boundaries.
Tone is too casual or too formal	Add the tone to the output spec. Show the tone in the examples.
Output occasionally refuses	Add a "flags" section explaining when to refuse and what to write instead.
Review takes too long	Add structure to the output so the reviewer can scan. Cut the output length range.

One last rule. An instruction that worked once is not a finished instruction. An instruction that worked on ten different inputs is a finished instruction. Aim for ten.

09 COLOPHON

The series and how to reach us.

Six field guides for UK SMEs. One Bristol-based company behind them.

The series

Six field guides for UK small and medium-sized businesses. Independent, plain-English, written by a practitioner.

01 · Where AI Actually Helps

Choose one worthwhile use case, scope a 30-day pilot and estimate the return before buying software.

02 · The 30-Day AI Pilot

A day-by-day guide to running the pilot you scoped in edition 01.

03 · Your Data, Your Rules

AI and privacy in plain English for UK businesses handling personal data.

04 · Writing AI Instructions That Actually Work

The six-part structure that turns vague prompts into reliable outputs.

05 · The AI Vendor Interrogation Kit

Eight questions to ask before signing a contract with any AI vendor.

06 · Your First AI Employee

A non-technical guide to autonomous AI agents for small businesses.

Colophon

PUBLICATION

Writing AI Instructions That Actually Work

July 2026 · Version 1.0

Author: Ted Milton, Expedient AI Ltd

CONTACT

Expedient AI Ltd

expedientai.co.uk

You can use this guide without contacting Expedient AI. If you would value an independent review of your data handling or AI controls, Expedient AI Ltd offers practical AI implementation support for UK SMEs.