

1 2 3 4 5 6 7

AI in 7 days

The ManuAI pocket guide

Understand AI in a minute. Use it the minute after.



Contents

How to use this guide	3
-----------------------	---

THE 7 DAYS	
Day 1: Meet AI	4

Day 2: Ask well	7
-----------------	---

Day 3: Know its limits	10
------------------------	----

Day 4: Stay safe	13
------------------	----

Day 5: Save time at work or school	16
------------------------------------	----

Day 6: Work like a manager	19
----------------------------	----

Day 7: Make your own AI habit	22
-------------------------------	----

Prompt recipes	25
----------------	----

What never to paste into AI	30
-----------------------------	----

Check AI's answer	31
-------------------	----

Glossary: 25 words	32
--------------------	----

Answers	35
---------	----

How to use this guide

Each day: three short lessons, one 10-minute task in your real life, and one question. About 15 minutes a day.

-
- 1 Read one day at a time. The lessons are short on purpose: each one is a single idea.
 - 2 Do the 10-minute task the same day, in a free chatbot such as ChatGPT, Claude, Gemini or Copilot. Each day ends on a page with room for your notes. Doing it is what makes it stick.
 - 3 Answer the day's question before you turn to the answers at the back.
 - 4 Blanks to fill in are written [LIKE THIS]. Replace them with your own details, and never paste private data (see the checklist).
 - 5 Every lesson is also online with a live toy to play with. The link sits under each lesson title.
-



Prefer to learn on screen? The same course, with progress saved and a certificate at the end:
ai.manusignal.com/7-days

01 DAY 1 Meet AI

3 lessons · 10-minute task · 1 question

GOAL Understand what a chatbot is and have your first useful conversation.

LESSON 1 · 1 MIN

ai.manusignal.com/lessons/what-is-ai

What is AI, really?

AI is software that learned patterns from examples instead of following rules a person wrote.

- Normal software follows instructions. A programmer writes: if the total is over 100, add free delivery.
- AI works differently. Instead of rules, it is shown millions of examples, and it finds the patterns by itself.
- Today, when people say AI they usually mean chatbots like ChatGPT, Claude or Gemini. These were trained on huge amounts of text, so they learned the patterns of language.
- AI is not a person and not magic. It is a very good pattern-matcher.

THINK OF IT LIKE THIS

Think of a child learning to recognise dogs. Nobody gives them a rulebook. They see many dogs, and one day they just know. AI learns the same way, but from millions of examples in a few weeks.

LESSON 3 · 1 MIN

ai.manusignal.com/lessons/what-a-chatbot-really-does

What a chatbot really does

A chatbot writes its answer one small piece at a time, each time picking a likely next word.

- Under the hood, ChatGPT, Claude and Gemini are large language models, or LLMs. Their core skill is simple: given some text, predict what comes next.
- Ask for a recipe for pho and it does not look up a recipe. It writes one, word by word, choosing words that usually follow in texts like this.
- This explains a lot. It is why answers sound fluent and confident.

THINK OF IT LIKE THIS

It is like the word suggestions on your phone keyboard, but trained on a library-sized amount of text and able to keep going for pages.

LESSON 4 · 1 MIN

ai.manusignal.com/lessons/tokens

Tokens: how AI reads

AI does not read words; it reads tokens, small chunks of text that are often parts of words.

- Before a model can read your message, the text is cut into tokens. A common English word like "the" is one token.
- Languages are not treated equally. The same sentence in Vietnamese, Hindi or Amharic often uses more tokens than in English, because the model saw less of those languages in training.
- Tokens matter for three reasons. Paid AI services charge per token. Every model has a limit on how many tokens it can handle at once. And odd token splits explain strange mistakes, like miscounting the letters in a word.

THINK OF IT LIKE THIS

Think of Lego. The model does not see a finished house; it sees the bricks. Common words are one big brick, rare words are built from several small ones.

Do it today

10-MINUTE TASK

Open a free chatbot (ChatGPT, Claude, Gemini or Copilot). Ask it to explain something you have always wondered about, then ask two follow-up questions. Notice how it builds on what you said.

YOUR NOTES

CHECK YOURSELF

What is a chatbot doing when it writes an answer?

A Looking up a stored answer

B Predicting likely next words, one piece at a time

C Asking a human

D Copying Wikipedia

Answer on page 35.

02

Ask well

DAY 2

3 lessons · 10-minute task · 1 question

GOAL Write prompts that get a good answer on the first try.

LESSON 7 · 2 MIN

ai.manusignal.com/lessons/the-four-part-prompt

The four-part prompt

Good answers come from prompts with four parts: a role, a task, some context, and a format.

- Most bad AI answers come from short, vague prompts. "Write an email" gives the model almost nothing to work with, so it guesses.
- Add four parts. Role: who should it act as? Task: what exactly should it do? Context: who is it for, and what does it need to know? Format: how long, what shape, what tone?
- You do not need all four every time. But when an answer disappoints, one of them is usually missing. Context is the one people forget most.

THINK OF IT LIKE THIS

It is like briefing a new colleague. "Do the report" fails. "You are our finance assistant. Write the monthly sales report for the Jakarta team, using these numbers, one page, bullet points" works.

LESSON 8 · 1 MIN

ai.manusignal.com/lessons/show-it-examples

Show it examples

One or two good examples in your prompt teach AI your style faster than a paragraph of instructions.

- Describing a style is hard. "Friendly but professional" means different things to different people.
- This is called few-shot prompting. You give a few examples of input and the output you want, then give the new input.
- It works well for product descriptions, social posts, replies to customers, sorting messages into categories, and anything where you already have a few good past samples.

- Use two or three examples that differ a little, so the model learns the pattern and not just one sentence. Remove any private details first.

THINK OF IT LIKE THIS

It is like showing a tailor a shirt you love instead of trying to describe the fit in words.

LESSON 10 · 1 MIN

ai.manusignal.com/lessons/let-it-interview-you

Let it interview you

When you do not know what to tell AI, ask it to ask you questions first.

- Writing a full prompt with all the context is hard. Often you do not know which details matter.
- The fix is simple: describe your goal in one line, then tell the AI to ask you the questions it needs before it starts. Answer them, and only then let it work.
- This turns a lazy prompt into a great one without much effort. It works very well for CVs, business plans, trip planning, difficult letters and lesson plans.
- Ask for a small number of questions at a time, like three to five. Otherwise you get a long list that feels like a form.

THINK OF IT LIKE THIS

It is like a good doctor. They do not prescribe after one sentence; they ask a few questions first.

Do it today

10-MINUTE TASK

Pick a real message you need to write this week. Write it as a four-part prompt (role, task, context, format), or ask the AI to interview you first. Edit the result and actually send it.

YOUR NOTES

CHECK YOURSELF

Which part of a prompt do people forget most often?

- A** Role

- B** Task

- C** Context

- D** Saying please

Answer on page 35.

03

DAY 3

Know its limits

3 lessons · 10-minute task · 1 question

GOAL Learn why AI gets things wrong and how to check it.

LESSON 5 · 2 MIN

ai.manusignal.com/lessons/why-ai-makes-things-up

Why AI makes things up

AI is built to produce a likely answer, not a true one, so sometimes it confidently invents facts.

- When a chatbot invents something, like a fake book title, a wrong date or a law that does not exist, it is called a hallucination.
- It happens because the model is predicting text that looks right. If it does not know the answer, it rarely says so.
- Hallucinations are most common with exact details: names, numbers, quotes, links, citations and recent events. They are less common with general explanations.
- Newer models hallucinate less, and many can search the web and show sources. But no model is free of it yet.

THINK OF IT LIKE THIS

It is like a student in an exam who did not study one topic but writes a confident paragraph anyway, because a blank page feels worse than a guess.

LESSON 14 · 2 MIN

ai.manusignal.com/lessons/check-what-ai-tells-you

Check what AI tells you

Use AI for a first draft of the truth, then check anything that matters before you act on it.

- AI can be wrong while sounding completely sure. For anything important, like health, money, law, safety or school work, check before you trust.
- A quick method: ask for sources, open them, and see if they really say what the AI claims. Fake or broken links are a warning sign.
- Cross-check with a second source: an official website, a second AI tool, or a person who knows. If they disagree, dig deeper.

- Be extra careful with numbers, dates, quotes and anything recent.

LESSON 6 · 2 MIN

ai.manusignal.com/lessons/context-window

The context window: AI's short-term memory

A chatbot can only see a limited amount of text at once, and in a very long chat the oldest parts can slip away.

- The context window is everything the model can see right now: your messages, its replies, any files you added, and hidden instructions. It is measured in tokens.
- Modern windows are large. Many models can hold a whole book.
- When a chat gets very long, the tool may drop or squash the oldest messages. That is why a chatbot can forget a rule you gave it an hour ago.

THINK OF IT LIKE THIS

It is like a whiteboard. You can write a lot, but when it fills up, someone wipes the top to make space. Anything wiped is gone unless you write it again.

Do it today

10-MINUTE TASK

Ask a chatbot three factual questions about your town, job or hobby. For each answer, ask for a source and check it. Write down which answers were right, wrong or unclear.

YOUR NOTES

CHECK YOURSELF

When should you check what AI tells you?

- | | |
|---|-----------------------------------|
| A Never, it is usually right | B Only for maths |
| C Whenever the answer matters:
health, money, law, school or work | D Only if it sounds unsure |

Answer on page 35.

04

Stay safe

DAY 4

3 lessons · 10-minute task · 1 question

GOAL Protect your data, your money and your family.

LESSON 13 · 1 MIN

ai.manusignal.com/lessons/what-not-to-paste

What not to paste into AI

Treat a chatbot like a public place: never paste anything you would not want a stranger to read.

- What you type into a chatbot is sent to the company's servers. Depending on the tool and your settings, it may be stored, reviewed by staff, or used to train future models.
- Never paste passwords, bank or card numbers, ID or passport numbers, medical records, or other people's private details. At work, do not paste customer data, contracts or secret plans unless your company has approved that tool.
- Before pasting, remove or replace sensitive bits: "Customer A" instead of a name, "[AMOUNT]" instead of the real number. The AI rarely needs the real details to help.
- Check the settings. Most tools let you turn off training on your chats, and work or school accounts often have stronger privacy rules.

THINK OF IT LIKE THIS

It is like talking on a phone in a busy train. You can discuss your plans, but you would not read out your PIN.

LESSON 15 · 2 MIN

ai.manusignal.com/lessons/spot-ai-scams

Spot AI-powered scams

Scammers now use AI to write perfect messages, so good grammar is no longer a sign that a message is real.

- Old scam emails were full of spelling mistakes. Now AI writes clean, friendly, personal messages in any language, at huge scale.
- Common tricks: a fake delivery text with a link, a bank warning asking you to log in, a job offer that pays too well, or a message from your boss asking you to buy gift cards. This is called phishing.

- Look for the pressure, not the grammar. Urgency ("in 30 minutes"), secrecy ("don't tell anyone"), and unusual payment (gift cards, crypto, a new bank account) are the real red flags.
- Never use the link or number in the message. Open the official app or call a number you already know.

LESSON 16 · 2 MIN

ai.manusignal.com/lessons/deepfakes

Deepfakes and voice clones

AI can now fake a face or a voice from a few seconds of real recording, so seeing and hearing is no longer proof.

- A deepfake is a fake video or image made with AI, often showing a real person saying or doing something they never did. Voice cloning copies how someone sounds from a short clip.
- Scammers use them to call parents with a crying "child" in trouble, to fake a CEO on a video call, or to show celebrities promoting fake investments.
- Visual clues like strange hands or blinking are getting rare. Better defence: check the story through a second channel.
- Agree a family safe word that only you know, for emergency calls.

THINK OF IT LIKE THIS

It is like a perfect forged signature. You cannot tell by looking, so you check with the person who supposedly signed it.

Do it today

10-MINUTE TASK

Open your chatbot's settings and decide whether it may train on your chats. Then agree a family safe word with someone close to you, for emergency calls asking for money.

YOUR NOTES

CHECK YOURSELF

Which is the biggest red flag in a message?

A Perfect grammar

B Urgent pressure to pay in an unusual way

C It uses your name

D It came by email

Answer on page 35.

05

DAY 5

Save time at work or school

3 lessons · 10-minute task · 1 question

GOAL Hand off the boring parts of writing and reading.

LESSON 19 · 1 MIN

ai.manusignal.com/lessons/email-and-documents

Emails and documents

Let AI write the first draft, then spend your time making it sound like you and getting the facts right.

- Writing is where AI saves most people the most time. A reply that takes you 15 minutes can take 2: you give the key points, it drafts, you edit.
- Give it the facts, the reader and the goal. "Reply to the supplier in Lagos: the delivery is 3 days late, we still want it, ask for a new date, keep it friendly."
- Watch for the AI voice: too long, too polite, too many fancy words. Ask it to be shorter and plainer, or give an example of how you write.
- Always read before sending. Check names, dates, numbers and promises.

LESSON 20 · 1 MIN

ai.manusignal.com/lessons/summarise-anything

Summarise anything

AI can turn a 40-page report into one page in seconds, if you tell it what you need the summary for.

- Most assistants accept long text, PDFs, web pages, even video transcripts. Summaries are one of the most reliable things AI does.
- A plain "summarise this" gives a generic result. Say who it is for and why: "Summarise for a busy manager who must decide whether to renew this contract."
- Ask for a structure: three key points, risks, numbers, and open questions. That makes summaries easy to scan and compare.
- Summaries can still miss or twist details. For decisions, ask it to quote the exact sentence it relied on, so you can check quickly.

Meetings and notes

AI can turn a messy meeting into clear decisions and a to-do list, as long as everyone knows it is listening.

- Video tools like Zoom, Teams and Google Meet now offer AI notes. You can also paste a transcript or your own rough notes into any chatbot.
- The most useful output is not a summary. It is three lists: decisions made, action items with an owner and a date, and questions still open.
- Always tell people when AI is recording or transcribing. In many countries it is the law, and everywhere it is good manners.
- AI notes can mishear names and numbers, or give an action to the wrong person. Have one human check the list before it is sent out.

Do it today

10-MINUTE TASK

Take one long document or a set of meeting notes. Ask AI for a summary for a specific reader, plus a list of actions with owners. Check one detail against the original.

YOUR NOTES

CHECK YOURSELF

How do you get a more useful summary?

- | | |
|--|---|
| A Ask for it to be as short as possible | B Say who it is for and what they need to decide |
| C Use bigger words | D Paste it twice |

Answer on page 35.

06

DAY 6

Work like a manager

3 lessons · 10-minute task · 1 question

GOAL Get better results by breaking jobs down and reviewing them.

LESSON 9 · 1 MIN

ai.manusignal.com/lessons/think-step-by-step

Ask it to think step by step

For problems with several steps, asking AI to work through them out loud gives more accurate answers.

- If you ask a tricky question and want only the answer, the model must jump straight to it. For maths, planning or logic, that jump often fails.
- Asking it to reason step by step is called chain-of-thought. The model writes out each step, and each step helps the next one.
- Many newer tools include reasoning models that do this thinking automatically, sometimes showing a "thinking" note. For those, you rarely need to ask.
- It is not needed for simple tasks like rewriting a paragraph. Use it for anything with numbers, conditions or several stages.

LESSON 11 · 2 MIN

ai.manusignal.com/lessons/prompt-chaining

Prompt chaining: one step at a time

Big jobs go better when you split them into small prompts and check each result before the next.

- Asking AI to "write my whole business plan" in one go usually gives a long, shallow answer. There is too much to do at once.
- Prompt chaining means breaking the job into steps. First, list the sections.
- You stay in control. If step two goes wrong, you fix it before it spoils step three.
- A good chain for most writing jobs: outline, then draft, then critique, then rewrite. It takes a few more minutes and gives a much better result.

THINK OF IT LIKE THIS

It is like cooking a meal. You do not throw every ingredient in the pot at once; you prepare, then cook, then taste and adjust.

Review AI like a manager

Treat AI like a fast, eager new hire: give clear briefs, check the work, and give specific feedback.

- A good manager does not accept the first draft or redo it all alone. They say what is wrong and ask for another version.
- Vague feedback like "make it better" gets vague changes. Specific feedback works: "The second paragraph repeats the first. Cut it. The price is wrong; it is 45 euros."
- Check four things every time: Is it true? Is it complete? Does it fit the reader? Would I sign my name on it?
- You can even ask AI to review itself first: "Act as a strict editor. List the five biggest problems with this draft." Then decide which to fix.

THINK OF IT LIKE THIS

AI is like a talented intern on day one. Very fast, sometimes brilliant, never tired, but it does not know your context and will not tell you when it is guessing.

Do it today

10-MINUTE TASK

Choose a bigger task, like a short plan or proposal. Do it as a chain: outline, draft, ask AI to critique the draft, then rewrite. Compare the final version with the first draft.

YOUR NOTES

CHECK YOURSELF

Which feedback gets the best second draft?

- | | |
|---|---------------------------------------|
| A "Make it better" | B "Try again" |
| C Specific changes: what to cut, fix and add | D No feedback, just regenerate |

Answer on page 35.

07 Make your own AI habit

DAY 7

3 lessons · 10-minute task · 1 question

GOAL See where AI is going and build a small weekly habit that keeps you improving.

LESSON 18 · 2 MIN

ai.manusignal.com/lessons/choosing-a-model

Choosing a model

There is no single best AI; pick the tool by the job, your budget and what your workplace allows.

- Popular chat assistants include ChatGPT, Claude, Gemini, Copilot, Meta AI, and others like DeepSeek, Mistral and Qwen. They leapfrog each other every few months, so rankings go out of date fast.
- Most offer a fast everyday model and a slower reasoning model for harder problems. Use fast for emails and quick questions, reasoning for maths, planning and analysis.
- Other things that matter: does it read images, PDFs and audio (multimodal)? Can it search the web? Does it work in your language? Is there an open-source model you can run on your own computer for privacy?
- At work, the first rule is policy. Use the tool your company approved, because it usually comes with better data protection.

LESSON 29 · 2 MIN

ai.manusignal.com/lessons/will-ai-take-my-job

Will AI take my job?

AI is more likely to change your tasks than remove your whole job, and people who use it well have the advantage.

- Jobs are bundles of tasks. AI is good at some: drafting, summarising, sorting, translating, basic analysis.
- Some roles will shrink, especially ones made mostly of routine writing or data entry. New tasks are also appearing: checking AI work, designing workflows, teaching others.
- The change is uneven across countries and industries, and nobody can predict it exactly. Be wary of anyone who is very sure, in either direction.

- The practical move is the same everywhere: find which of your tasks AI can help with, get good at it, and spend the saved time on the human parts.

LESSON 30 · 1 MIN

ai.manusignal.com/lessons/how-to-keep-up

How to keep up without burning out

You do not need to follow every AI headline; a small weekly habit beats chasing the news.

- New models and tools appear almost every week. Trying to follow it all is exhausting, and most news will not change how you work.
- Focus on the basics that last: how models predict, why they hallucinate, how to write a good prompt, how to stay safe. These lessons stay true even as tools change.
- Then build a tiny habit. Once a week, take one real task and try doing it with AI.
- Pick one or two trusted sources for news, not ten. And when a new tool claims to change everything, wait a month and see if people still use it.

THINK OF IT LIKE THIS

It is like fitness. You do not need every new workout trend; you need a simple routine you actually keep.

Do it today

10-MINUTE TASK

Make your own AI habit: list your weekly tasks and pick the one where AI helped most this week. Save your best prompt from this course as a template, and put a 15-minute "try one task with AI" slot in your calendar every week.

YOUR NOTES

CHECK YOURSELF

What keeps your AI skills growing after this course?

- | | |
|--|-------------------------------------|
| A Reading every AI headline | B Trying a new app every day |
| C Using AI on one real task each week and noting what works | D Waiting for a perfect AI |

Answer on page 35.

Prompt recipes

Copy-ready prompts. Each follows role + task + context + format, and the blanks to fill in are written as [UPPERCASE BRACKETS].

01 Rewrite it clearer

WRITING

You are a careful editor who values plain language.

Rewrite the text below so it is clearer and easier to read, without changing its meaning or dropping any facts.

It is for [AUDIENCE], who will read it [WHERE OR HOW, E.G. ON A PHONE].

Return the rewritten text first, then a short bullet list of the main changes you made.

Text:

[PASTE YOUR TEXT]

TIP Read the change list: it shows you habits to fix in your own writing next time.

02 Reply to an email

WRITING

You are my assistant, skilled at writing friendly, professional emails.

Draft a reply to the email below. I want to [WHAT YOU WANT TO SAY OR ACHIEVE].

Context: my relationship to the sender is [E.G. CLIENT, MANAGER, NEIGHBOUR], and the tone should be [E.G. WARM, FIRM, BRIEF].

Keep it under [NUMBER] words, with a clear subject line if one is needed.

Email I received:

[PASTE THE EMAIL]

TIP Say what you want the reply to achieve; that matters more than the wording.

03 Make it shorter

WRITING

You are an editor who cuts words, not meaning.
Shorten the text below to about [TARGET LENGTH, E.G. 100 WORDS OR HALF].
Keep these points no matter what: [MUST-KEEP POINTS, OR "THE MAIN ASK"].
Return only the shortened text.

Text:
[PASTE YOUR TEXT]

TIP Naming the must-keep points stops important details being cut.

04 Pros and cons

THINKING

You are a balanced advisor who helps people think, not decide for them.
I am deciding whether to [DECISION].
My situation: [KEY FACTS, E.G. BUDGET, TIMING, WHO IS AFFECTED].
What matters most to me: [YOUR PRIORITIES].
List the main pros and cons in two columns, mark which ones are most important given my priorities, and finish with three questions I should answer before deciding.

TIP The closing questions are often more useful than the list itself.

05 Break down a big task

THINKING

You are a practical project planner.
I need to [BIG GOAL] by [DEADLINE]. I can spend about [TIME PER WEEK] on it.
What I have already: [RESOURCES, HELP, PROGRESS SO FAR].
Break this into small steps of no more than [E.G. 1 HOUR] each, grouped into stages, in the order I should do them. Mark the first step I can start today.

TIP Small first steps beat a perfect plan; start the marked one right away.

06 Pre-mortem

THINKING

You are an experienced risk advisor.

Imagine it is [TIME FROM NOW, E.G. SIX MONTHS] later and my plan has failed.

The plan: [DESCRIBE THE PLAN]

Write the most likely story of how it failed. Then list the top five causes as a table with columns: cause, early warning sign, what I can do now to prevent it.

TIP Imagining failure in advance makes risks easier to see than asking 'what could go wrong'.

07 Explain it like I'm new

LEARNING

You are a patient teacher who explains without jargon.

Explain [TOPIC] to me. I am completely new to it; I know [WHAT YOU ALREADY KNOW, OR "NOTHING"].

Use one everyday analogy, keep it under 200 words, and define any technical word the first time you use it. End by asking me one question to check I understood.

TIP Answer its check question honestly; it will adjust the next explanation to fit you.

08 Quiz me

LEARNING

You are a friendly tutor.

Quiz me on [TOPIC OR PASTE YOUR NOTES] at [LEVEL, E.G. BEGINNER, EXAM LEVEL].

Ask one question at a time and wait for my answer. After each answer, tell me if I am right, explain briefly, then ask the next one. Make them gradually harder. After [NUMBER] questions, summarise what I should review.

TIP Being tested beats re-reading: getting answers wrong is part of how you learn.

09 Meeting summary

WORK

You are an efficient meeting secretary.

Summarise the meeting notes or transcript below for [AUDIENCE, E.G. PEOPLE WHO MISSED IT].

Format: a three-line overview, then headings for decisions made, open questions, and next steps (who, what, by when). Only include what is in the notes; write "unclear" rather than guessing.

Notes:

[PASTE NOTES OR TRANSCRIPT]

TIP Check names and dates before sending; transcripts often mix up speakers.

10 Pull out action items

WORK

You are a detail-focused project coordinator.

From the text below, extract every action item.

Return a table with columns: task, owner, deadline, source quote. If the owner or deadline is not stated, write "not stated".

Text:

[PASTE EMAIL THREAD, NOTES OR CHAT]

TIP The source quote column lets you check each item quickly.

11 Prioritise my to-do list

WORK

You are a calm, practical productivity coach.

Here is my to-do list: [PASTE LIST].

Context: my main goal this week is [GOAL], and I have about [HOURS] of free time.

Fixed deadlines: [DEADLINES].

Sort the tasks into: do today, schedule this week, delegate, and drop. Give a one-line reason for each, then suggest the order for today.

TIP Be honest about your available hours, or the plan will not survive the day.

12 Explain a confusing letter

LIFE

You are a helpful adviser who explains official documents in plain language.

Explain the letter or document below. I live in [COUNTRY].

Tell me: what it is about in two sentences, what I need to do and by when, any costs or risks, and what questions I should ask the sender.

Document (personal details removed):

[PASTE TEXT]

TIP Remove ID numbers and account details first, and confirm key deadlines with the sender.

13 Prepare for a hard conversation

LIFE

You are a calm communication coach.

I need to talk to [WHO, E.G. MY FLATMATE] about [ISSUE].

What I want from the conversation: [OUTCOME]. What I am worried about: [WORRIES].

Help me plan: a gentle opening line, the main points in my own voice, two ways they might react and how I could respond, and a fair compromise I could offer.

TIP Practise by asking it to play the other person, then try your opening out loud.

14 Reply to a customer

BUSINESS

You are a customer service lead known for calm, helpful replies.

Write a reply to the customer message below. Our business: [WHAT YOU SELL].

What we can offer: [E.G. REFUND, REPLACEMENT, APOLOGY ONLY]. Our policy: [RELEVANT POLICY].

Acknowledge their issue, explain clearly what happens next, and keep it under [NUMBER] words. Use a [TONE] tone.

Customer message:

[PASTE MESSAGE]

TIP Only promise what you can deliver; edit any offer the AI adds on its own.

What never to paste into AI

Treat a chatbot like a public place: never paste anything you would not want a stranger to read.

What you type into a chatbot is sent to the company's servers. Depending on the tool and your settings, it may be stored, reviewed by staff, or used to train future models.

NEVER PASTE

- ☒ Passwords
 - ☒ Bank or card numbers
 - ☒ ID or passport numbers
 - ☒ Medical records
 - ☒ Other people's private details
 - ☒ At work: customer data, contracts or secret plans, unless your company has approved that tool
-

BEFORE YOU PASTE

- ☐ Before pasting, remove or replace sensitive bits: "Customer A" instead of a name, "[AMOUNT]" instead of the real number. The AI rarely needs the real details to help.
 - ☐ Check the settings. Most tools let you turn off training on your chats, and work or school accounts often have stronger privacy rules.
-

Check AI's answer

Use AI for a first draft of the truth, then check anything that matters before you act on it.

CHECK BEFORE YOU ACT WHEN IT IS ABOUT

Health · Money · Law · Safety · School work

EVERY TIME

- ☐ A quick method: ask for sources, open them, and see if they really say what the AI claims. Fake or broken links are a warning sign.
 - ☐ A source that does not support the claim is no source. Treat the claim as unchecked until a reliable source confirms it.
 - ☐ Cross-check with a second source: an official website, a second AI tool, or a person who knows. If they disagree, dig deeper.
 - ☐ Hallucinations are most common with exact details: names, numbers, quotes, links, citations and recent events. They are less common with general explanations.
-

ASK IT TO SHOW ITS WORK

For your last answer, list each factual claim, how confident you are in it (high, medium, low), and a source I can check for it. Say clearly if you cannot find a source.

BEFORE YOU USE IT, ASK

- | | |
|--|--|
| ☐ Is it true? | ☐ Is it complete? |
| ☐ Does it fit the reader? | ☐ Would I sign my name on it? |
-

Glossary: 25 words

The words you will meet most often, in plain English. All of them, and more, at ai.manusignal.com/glossary.

Agent

An AI system that takes several steps on its own, such as searching, clicking or using apps, to finish a goal you set.

AI · Artificial Intelligence

Computer software that learns patterns from examples so it can recognise, predict or produce things, instead of following only hand-written rules.

Automation bias

Our habit of trusting a machine's answer too readily, even when our own judgement or other evidence says otherwise.

Bias

When an AI treats some people or ideas unfairly because of skewed patterns in the data it learned from.

Chain of thought

Asking an AI to work through a problem in steps before answering, which often makes its answers more accurate.

Chatbot

A program you talk to in everyday language by typing or speaking, which replies in the same way.

Context window

The maximum amount of text an AI can look at in one go. Anything beyond it is forgotten for that conversation.

Data privacy

Keeping personal or confidential information under control, including deciding what you type into AI tools and where it is stored.

Deepfake

A fake video, photo or audio clip made with AI to show a real person saying or doing something they never did.

Few-shot prompting

Including a few examples of what you want in your prompt so the AI copies their style and format.

Generative AI

AI that produces new text, images, audio, video or code, rather than only sorting or scoring existing things.

Hallucination

When an AI states something false with confidence, such as a made-up fact, quote or source.

Human in the loop

Designing a process so a person reviews or approves the AI's work before it has real consequences.

Knowledge cutoff

The date after which a model learned nothing new. Unless it can search the web, it may not know recent events.

LLM · Large Language Model

An AI trained on huge amounts of text to predict the next word, which lets it write, summarise, translate and chat.

Machine learning

Teaching computers by showing them many examples so they find the patterns themselves, instead of being given every rule.

Model

The trained file of learned patterns that does the actual work inside an AI product. One model can power many apps.

Next-word prediction

The core job of a language model: guessing the most likely next piece of text, over and over, to build a reply.

Phishing

Fake messages that pretend to be someone you trust to trick you into clicking a link, paying or sharing passwords.

Prompt

The instruction or question you give an AI. Clearer prompts with context usually get much better answers.

Prompt chaining

Breaking a big job into a sequence of smaller prompts, feeding each answer into the next step.

Sycophancy

An AI's tendency to agree with you or flatter your ideas instead of telling you where you are wrong.

Token

A chunk of text, often part of a word, that AI reads and writes in. Limits and prices are counted in tokens.

Training data

The examples, such as text, images or recordings, that a model learns its patterns from.

Voice cloning

Using AI to copy someone's voice from a short recording so it can say anything. Common in phone scams.

Answers

The answers to each day's question, with why.

DAY 1 What is a chatbot doing when it writes an answer?

B Predicting likely next words, one piece at a time

Chatbots generate text by predicting what comes next, again and again.

DAY 2 Which part of a prompt do people forget most often?

C Context

Context, like who it is for and the key facts, is what turns a generic answer into a useful one.

DAY 3 When should you check what AI tells you?

C Whenever the answer matters: health, money, law, school or work

AI sounds confident even when wrong, so check anything you will act on.

DAY 4 Which is the biggest red flag in a message?

B Urgent pressure to pay in an unusual way

AI makes scam messages look polished. Urgency plus strange payment methods is the real warning sign.

DAY 5 How do you get a more useful summary?

B Say who it is for and what they need to decide

A clear reader and goal tells the AI what to keep and what to leave out.

DAY 6 Which feedback gets the best second draft?

C Specific changes: what to cut, fix and add

Specific feedback works with AI just as it does with people.

DAY 7 What keeps your AI skills growing after this course?

C Using AI on one real task each week and noting what works

A small, regular habit on real work builds more skill than chasing news.

Keep going online

Everything in this guide lives at ai.manusignal.com, free and with no account, along with the live toys, quizzes and case files.

THE 7-DAY COURSE, WITH YOUR PROGRESS SAVED

ai.manusignal.com/7-days

ALL THE LESSONS

ai.manusignal.com/lessons


YOUR CERTIFICATE

ai.manusignal.com/certificate

Finish all seven days to unlock your certificate.

TIẾNG VIỆT

ai.manusignal.com/vi



Know someone who should read this? Send them the PDF. It is free to share.