


◆ — ENGLISH EDITION — ◆

AI MASTERY

From No Clue to Master

The complete beginner-to-expert course on using Artificial Intelligence — no technical background required.

- 
- 01 No Clue
 - 02 Novice
 - 03 Advanced
 - 04 Pro
 - 05 Master

Welcome

Most people use Artificial Intelligence the way they'd use a magic 8-ball: type something vague, get something mediocre, and conclude that "AI is overrated." The truth is that AI is a **skill** — and like every skill, there is a shortcut: learning it in the right order.

This course takes you through **5 phases**, from absolute zero to genuine mastery. Each phase builds directly on the one before it. By the end you will not just "chat with a bot" — you will delegate real work to AI, fix it when it goes wrong, and build systems that work while you sleep.

You do not need any technical background. If you can write a message, you can do everything in this book.

How to Use This Course

- **One phase at a time.** Do not skip ahead. Each phase assumes you have truly absorbed the previous one.
- **Do the exercises as you read.** Open a free AI tool (Phase 1 shows you how) and try everything immediately. Reading about AI without using it is like reading about swimming.
- **Take the Standard Test at the end of each phase.** Ten questions, and the pass mark is 10/10 — not to be harsh, but because every question covers something you will genuinely need later. The Answer Keys are at the back of the book.
- **Come back to it.** This book is designed as a reference. The frameworks in Phases 3–5 are worth re-reading once you have a few weeks of practice.

© 2026 Paper Plane Workshop. Free edition — you are welcome to share the download link (paperplaneworkshop.com), but this book may not be resold or modified. The information in this book is provided for educational purposes; AI tools change quickly, and specific product features may differ by the time you read this.

Table of Contents

Welcome 2

 How to Use This Course 2

PHASE 1 — No Clue 5

 Lesson 1.1 — What AI Actually Is (in Plain Language) 5

 Lesson 1.2 — Meet the Big Tools 5

 Lesson 1.3 — Your First Conversation 6

 Lesson 1.4 — What AI Can and Can't Do 6

 Lesson 1.5 — Staying Safe from Day One 7

PHASE 2 — Novice 11

 Lesson 2.1 — The Anatomy of a Good Prompt 11

 Lesson 2.2 — Everyday Wins 11

 Lesson 2.3 — Iterate: The Novice's Superpower 12

 Lesson 2.4 — AI in Your Pocket 12

 Lesson 2.5 — Learning Anything with AI 12

PHASE 3 — Advanced 16

 Lesson 3.1 — The R-T-C-F Framework 16

 Lesson 3.2 — Working with Documents 16

 Lesson 3.3 — Your Prompt Library 17

 Lesson 3.4 — Creating Beyond Text 17

 Lesson 3.5 — Multi-Step Projects 18

PHASE 4 — Pro 22

 Lesson 4.1 — Diagnosing Bad Output 22

 Lesson 4.2 — Expert Techniques 22

 Lesson 4.3 — Controlling Hallucination 23

 Lesson 4.4 — Choosing Models & Managing Cost 23

 Lesson 4.5 — Automation: AI That Runs Without You 24

PHASE 5 — Master 27

Lesson 5.1 — Agents: AI That Takes Action	27
Lesson 5.2 — Your Own Knowledge Base	27
Lesson 5.3 — APIs & Building Without Code	28
Lesson 5.4 — Designing AI Workflows for Business	28
Lesson 5.5 — Staying Ahead & Turning Skill Into Value	29
Answer Keys	33
You Made It — Now What?	36

No Clue

Absolute beginner fundamentals

This phase assumes you know **nothing** — and that's a great place to start, because most confusion about AI comes from half-learned jargon. In the next five lessons you will understand what AI actually is, meet the main tools, have your first real conversation with one, and learn the safety rules that even experienced users break.

Lesson 1.1 — What AI Actually Is (in Plain Language)

The AI everyone is talking about is a type of software called a **Large Language Model (LLM)**. Despite the intimidating name, the core idea is simple: an LLM has read a colossal amount of text — books, articles, websites, code — and learned the **patterns** in how humans write and reason. When you send it a message, it uses those patterns to predict the most useful continuation, one word at a time.

That single trick turns out to be astonishingly powerful. Because human knowledge is written down in text, a system that deeply understands text can **explain, summarize, translate, write, plan, and reason** across almost every topic.

Three terms you will see constantly:

- **AI (Artificial Intelligence)** — the broad umbrella term for software that performs tasks that normally need human intelligence.
- **LLM (Large Language Model)** — the specific technology behind ChatGPT, Claude, and Gemini. It works with language.
- **Chatbot / Assistant** — the app you actually talk to. The chatbot is the car; the LLM is the engine.

One mental model to keep forever: an LLM is **not a database** that looks up answers. It is closer to an extremely well-read colleague answering from memory — usually brilliant, occasionally confidently wrong. That difference drives everything else in this course.

Lesson 1.2 — Meet the Big Tools

You do not need to try every AI product — the market leaders are similar enough that skills transfer between them. The main names to know:

- **ChatGPT** (by OpenAI) — the most famous one; strong all-rounder with a generous free tier.

- **Claude** (by Anthropic) — known for careful, high-quality writing and handling long documents.
- **Gemini** (by Google) — tightly integrated with Google Search, Gmail, and Docs.
- **Copilot** (by Microsoft) — built into Windows, Edge, and Microsoft Office.

All of them offer a **free tier** that is more than enough for this course. Paid plans (usually around \$20/month) buy you a smarter model, more messages, and extra features — worth it later, unnecessary today.

Each tool exists as a **website** you log into and a **mobile app**. Pick any one of them, create a free account with your email, and keep it open while you read the rest of this book. Everything you learn here works in all of them.

Lesson 1.3 — Your First Conversation

Open your chosen tool. You will see a text box. This is a **conversation**, not a search engine — and that changes how you should use it:

- **Write in full sentences**, like a message to a smart assistant: "Explain how mortgages work like I'm 15" beats "mortgage explanation".
- **It remembers the conversation**. You can say "make it shorter" or "now give me an example" and it knows what "it" means. Each chat is a running thread.
- **A new chat starts fresh**. Open a new conversation for a new topic — the AI does not carry memory between chats (unless the app has an explicit memory feature).

Try this right now — type: **"I'm a complete beginner with AI. Ask me 3 questions about my life or work, then suggest 5 ways you could genuinely help me."** Notice something important: you just instructed the AI to interview *you*. Steering the conversation like this is your first real prompt skill.

Do not worry about phrasing things "correctly." There are no magic words, no commands to memorize, and you cannot break anything. Typos are fine. The AI's job is to understand you.

Lesson 1.4 — What AI Can and Can't Do

AI is dramatically better than most beginners expect at some things — and worse at others. Knowing the difference is the real fundamental skill of Phase 1.

Where AI is strong: writing and rewriting anything (emails, letters, posts, CVs), summarizing long text, explaining difficult topics at any level, translating, brainstorming ideas, planning (trips, projects, meals), and acting as a patient tutor on almost any subject.

Where AI is weak:

- **Hallucination** — the number one weakness. When an LLM doesn't know something, it may generate a confident, plausible, completely false answer: fake statistics, invented book titles, wrong dates. It is not lying; it is pattern-completing.
- **Fresh events** — models are trained up to a cutoff date. Many tools can search the web to compensate, but only when that feature is on.
- **Your private context** — it does not know your company, your inbox, or your life unless you tell it in the conversation.
- **Precise arithmetic on big numbers** — it predicts text, it is not a calculator (though many tools can call one).

This gives us the **Golden Rule of AI**, which you should treat as law from today: *use AI freely for drafts, ideas, and explanations — but verify any fact, number, name, or quote before you rely on it or repeat it.*

Lesson 1.5 — Staying Safe from Day One

Five safety habits, established now, will protect you for your entire AI journey:

- **Never paste secrets.** No passwords, card numbers, ID numbers, or other people's private data. Treat a chat like a postcard, not a vault.
- **Assume inputs may train future models** on free consumer tiers unless settings say otherwise. Most tools have a setting to opt out of training — turn it off in Settings on day one.
- **Verify before you act.** Medical, legal, and financial answers are starting points for research, not professional advice.
- **Watch for scams.** The AI boom created fake apps and "AI investment" fraud. Download only official apps from the companies named in Lesson 2.
- **Disclose when it matters.** If AI wrote most of something important that others assume is fully yours, be honest about it. Reputation outlasts convenience.

None of this should scare you off — these are the same common-sense rules as for email or online banking. Set them once and they become invisible.

Key Takeaways

- ✓ An LLM predicts text based on patterns learned from vast reading — it's a brilliant colleague answering from memory, **not** a database.
- ✓ ChatGPT, Claude, Gemini, and Copilot are the main tools; all have free tiers, and skills transfer between them.
- ✓ Talk to AI in full sentences, in a running conversation — you can always refine with follow-ups.
- ✓ AI excels at writing, summarizing, explaining, and planning; it can **hallucinate** confident false facts.
- ✓ Golden Rule: draft with AI freely, but **verify** every important fact. Never paste secrets.

— PHASE 1 · STANDARD TEST —

Answer all 10 questions, then check yourself against the Answer Key at the back of this book. You need 10/10 to consider the phase passed. If you miss a question, re-read the matching lesson and retake the test — that repetition is exactly how the knowledge sticks.

1. What does an LLM fundamentally do?

- A) Looks up answers in a built-in database of facts
- B) Predicts the most useful continuation of text based on learned patterns
- C) Searches the internet for every answer
- D) Runs pre-written scripts created by programmers

2. Which of these is the relationship between a chatbot and an LLM?

- A) They are competing technologies
- B) The LLM is the app; the chatbot is the engine
- C) The chatbot is the app you talk to; the LLM is the engine powering it
- D) A chatbot is just a smaller LLM

3. What is 'hallucination' in AI?

- A) The AI refusing to answer a question
- B) The AI generating confident but false information
- C) A visual glitch in the app
- D) The AI remembering a previous conversation

4. You asked the AI for a summary and want it shorter. What's the best next step?

- A) Open a new chat and start over
- B) Reply 'make it shorter' in the same conversation
- C) Report the answer as wrong
- D) Retype the whole original request with the word 'short'

5. Which task is an LLM WEAKEST at, out of the box?

- A) Rewriting an email in a friendlier tone
- B) Explaining a physics concept simply
- C) Reporting this morning's breaking news without web access
- D) Brainstorming names for a bakery

6. What is the Golden Rule of AI from this phase?

- A) Always use the paid version for accuracy
- B) Use AI for drafts and ideas, but verify important facts before relying on them
- C) Never use AI for anything factual
- D) Trust the AI unless it sounds unsure

7. Which of these is SAFE to paste into an AI chat?

- A) Your online banking password
- B) A friend's ID number to fill in a form
- C) A news article you want summarized
- D) Your full card number to 'check if it's valid'

8. Why does a brand-new chat not remember your previous conversation?

- A) It's a bug in most AI tools
- B) Each chat is a separate thread; memory doesn't carry over by default
- C) The AI deletes data every hour for legal reasons
- D) It does remember — it just pretends not to

9. Roughly what does a typical AI paid plan (~\$20/month) get you?

- A) It's the only way to use AI at all
- B) A smarter model, higher limits, and extra features
- C) A guarantee that answers are always correct
- D) Legal permission to use answers commercially

10. Which first message best applies this phase's advice?

- A) 'weather'
- B) 'Explain how car leasing works like I'm 15, then give me the 3 most common mistakes people make.'
- C) 'Do you work?'
- D) 'asdf test 123'

Novice

Basic practical applications

You know what AI is and you've had your first conversation. Phase 2 is where AI starts **saving you real time**. The theme of this phase is simple: better instructions produce better results. You'll learn the anatomy of a good prompt, apply it to everyday tasks, and pick up the habit that separates novices from beginners — **iterating** instead of accepting the first answer.

Lesson 2.1 — The Anatomy of a Good Prompt

A **prompt** is simply the message you send to the AI. The quality gap between a lazy prompt and a good one is enormous — often the difference between a useless answer and a perfect one. Three ingredients fix 80% of bad prompts:

- **Be specific about the task.** Not "write an email" but "write a short, polite email rescheduling Friday's meeting to next Tuesday."
- **Give context.** Who are you? Who is it for? What's the situation? "I'm emailing my landlord, we're on friendly terms" changes everything about the tone.
- **State the format you want.** A table, a bullet list, 100 words maximum, three options to choose from — if you don't say, the AI guesses.

Compare: *"give me dinner ideas"* versus *"Give me 5 dinner ideas for two adults, under 30 minutes, vegetarian, high-protein, using rice or pasta as a base. Format: a table with dish name, cook time, and main ingredients."* Same effort to type — wildly different output.

You don't need perfect prompts. You need **complete** ones: task + context + format. Make that your checklist for every request that matters.

Lesson 2.2 — Everyday Wins

The fastest way to build the AI habit is to hand over tasks you already do every week. Proven starters:

- **Messages & email** — drafting, softening a tense reply, translating tone ("make this sound professional but warm").
- **Summarizing** — paste any long text and ask for 5 bullet points, or "explain what this contract clause means in plain language."
- **Translation** — modern LLMs translate naturally, keeping tone and idioms, and you can ask follow-up questions about nuances.

- **Planning** — trips ("3-day Rome itinerary, walking-focused, budget-friendly"), meal plans with shopping lists, event checklists.
- **Decisions** — "Compare these two phones for a non-technical parent. Table, then a recommendation."

Pick **three** of these and do them today with real tasks from your life. The point of this phase is not knowledge — it's converting knowledge into reflex.

Lesson 2.3 — Iterate: The Novice's Superpower

Beginners treat the first answer as final. Novices treat it as a **first draft**. The conversation is the tool — you refine output the same way you'd direct a helpful assistant:

- "Make it shorter / friendlier / more formal."
- "Give me 3 alternative versions."
- "That's too generic — mention that we've been customers for five years."
- "Now turn it into a bullet-point summary I can paste into chat."

Two more refinement moves worth knowing: **regenerate** (ask it to try again — you often get a noticeably different take) and **ask why** ("why did you choose this structure?") which frequently reveals options you didn't know you had.

Rule of thumb: for anything that matters, expect **2–4 rounds** of refinement. That's not the AI failing — that's the workflow working.

Lesson 2.4 — AI in Your Pocket

The mobile apps of ChatGPT, Claude, Gemini, and Copilot unlock three modes that most beginners never touch:

- **Voice mode** — hold a natural spoken conversation. Perfect while walking or driving: practice a language, rehearse a difficult conversation, get a topic explained hands-free.
- **Camera / photos** — snap a picture and ask about it: a plant to identify, a menu in another language, a confusing letter from a government office, an error message on a screen.
- **Sharing into the app** — most phones let you share a webpage, PDF, or email straight to the AI app and ask questions about it.

Photo-plus-question is the single most underrated novice skill: **"What is this and what should I do about it?"** works on rashes on a plant, warning lights on a dashboard, and paragraphs of legal jargon alike. (Golden Rule still applies — verify before acting on anything important.)

Lesson 2.5 — Learning Anything with AI

AI may be the best personal tutor ever created: infinitely patient, available at 3 a.m., and never judgmental about "dumb" questions. Techniques that make it click:

- **Set the level:** "Explain quantum computing like I'm 12" or "...assume I know high-school physics."
- **Ask for analogies:** "Give me an analogy from cooking" — analogies anchor new ideas to things you know.
- **Flip the direction:** "Ask me 5 questions to test whether I understood this, one at a time, and correct my answers."
- **Request a learning plan:** "Build me a 2-week plan to learn Excel basics, 30 minutes a day, with a small exercise each day."

The flip-the-direction trick matters beyond studying: an AI that **quizzes you** converts passive reading into active recall — the same principle this book uses with its phase tests. From now on, whenever you finish learning something important, ask the AI to test you on it.

Key Takeaways

- ✓ A complete prompt has three parts: **specific task + context + desired format**.
- ✓ Hand AI your recurring weekly tasks first — email, summaries, translation, planning — to build the reflex.
- ✓ The first answer is a draft. **Iterate 2–4 times:** shorter, friendlier, three versions, add missing detail.
- ✓ On mobile, use **voice mode** and **photo + question** — 'what is this and what should I do?'
- ✓ As a tutor, AI works best when you set the level, ask for analogies, and make it **quiz you**.

— PHASE 2 · STANDARD TEST —

Answer all 10 questions, then check yourself against the Answer Key at the back of this book. You need 10/10 to consider the phase passed. If you miss a question, re-read the matching lesson and retake the test — that repetition is exactly how the knowledge sticks.

1. What are the three ingredients of a complete prompt?

- A) Politeness, length, keywords
- B) Specific task, context, desired format
- C) Question, answer, follow-up
- D) Greeting, request, thank-you

2. Which prompt is MOST likely to get a useful answer first try?

- A) 'food ideas'
- B) 'Write something about meetings'
- C) 'Write a polite 3-sentence email moving Friday's team meeting to Tuesday 10:00, and suggest they reply with conflicts.'
- D) 'Can you help me with email?'

3. The AI's first draft of your email is too stiff. Best novice move?

- A) Give up and write it yourself
- B) Reply: 'Make it warmer and more casual, keep it under 100 words'
- C) Start a new chat and hope for better luck
- D) Send it anyway

4. Why is 'ask me questions to test my understanding' such a powerful learning prompt?

- A) It makes the AI respond faster
- B) It converts passive reading into active recall
- C) It shortens the AI's answers
- D) It avoids hallucinations completely

5. What can you do with a photo in a mobile AI app?

- A) Nothing — AI apps are text-only
- B) Only apply filters to it
- C) Ask questions about it, like translating a menu or explaining a document
- D) Only store it in the chat history

6. How many refinement rounds should you EXPECT for output that matters?

- A) Zero — good prompts get perfect first answers
- B) 2–4 rounds
- C) At least 20 rounds
- D) Exactly one regenerate

7. You want dinner ideas in an easy-to-compare layout. What should your prompt include?

- A) A request for a specific format, like a table with columns
- B) The word 'please' at least twice
- C) A demand for only one single answer
- D) Nothing — layout can't be controlled

8. Which is a genuinely useful voice-mode scenario?

- A) Practicing a language conversation while walking
- B) Sending passwords more securely by voice
- C) Making the AI's answers more accurate
- D) Unlocking secret models

9. What does 'Explain it like I'm 12' actually control?

- A) The AI's speed
- B) The complexity level of the explanation
- C) The truthfulness of the answer
- D) The language it replies in

10. Why does adding context like 'I'm writing to my landlord and we're on friendly terms' improve output?

- A) It makes the message longer, which is always better
- B) It lets the AI match tone and content to the real situation
- C) Landlords are a special AI category
- D) It doesn't — context is ignored

Advanced

Core system building & deeper concepts

Phases 1 and 2 made you a competent user. Phase 3 makes you a **systematic** one. The difference: a novice writes a good prompt when they need one; an advanced user has a **framework** for constructing prompts, a **library** of tested templates, and knows how to feed real documents and multi-step projects into AI. This is the phase where AI stops being a toy and becomes infrastructure.

Lesson 3.1 — The R-T-C-F Framework

Everything from Phase 2 condenses into one professional framework: **R-T-C-F** — Role, Task, Context, Format. For any important prompt, supply all four:

- **R — Role:** who should the AI be? "You are an experienced employment lawyer" / "You are a friendly fitness coach for beginners." A role activates the right vocabulary, priorities, and depth.
- **T — Task:** the specific action, with a verb. Analyze, rewrite, compare, generate 10 options, critique.
- **C — Context:** the situation, the audience, what you've tried, constraints, and any relevant text pasted in.
- **F — Format:** structure (table, numbered steps, JSON), length, tone, and language.

A full R-T-C-F prompt: "**You are a senior recruiter (R). Rewrite my CV summary below to target a junior data-analyst role (T). Context:** I'm switching from retail management, my Excel and SQL are self-taught, employers worry about the career change (C). **Format:** 3 versions, each max 60 words, confident but not arrogant (F)."

R-T-C-F is not a ritual — skip parts when stakes are low. But when output quality matters, running through the four letters takes twenty seconds and reliably transforms results. It is the single highest-value habit in this book.

Lesson 3.2 — Working with Documents

Modern AI tools accept **file uploads** (PDF, Word, spreadsheets, images) and huge pasted texts. This unlocks a new class of work:

- **Summarize & extract:** "Summarize this 40-page report in one page, then list every mentioned deadline and amount in a table."

- **Explain:** "Explain section 4 of this contract in plain language. What are the risks for me as the tenant?"
- **Compare:** upload two offers, two CVs, two insurance policies — "compare in a table, then recommend."
- **Transform:** turn meeting notes into an action list, a spec into an email, a spreadsheet into a written summary.

Two professional habits: first, **tell the AI what you'll use the output for** — a summary "to decide whether to attend this conference" differs from one "to brief my CEO in 30 seconds." Second, for long documents, ask questions **in stages**: overview first, then drill into sections. You get better answers than with one giant request.

Privacy checkpoint from Phase 1 still applies: anonymize sensitive documents (swap names, remove account numbers) before uploading when the content is confidential.

Lesson 3.3 — Your Prompt Library

Advanced users never write the same good prompt twice. When a prompt works well, **save it** — a notes app or document is enough. That collection is your **prompt library**, and it compounds like savings.

Structure each entry as: a name, the prompt template with **placeholders in [brackets]**, and a note on when to use it. Example: *"Email Defuser — You are a calm communication expert. Rewrite my draft below to de-escalate while keeping my key request: [DRAFT]. Keep it under [N] words."*

Then go one level up with **custom instructions** (ChatGPT) / **preferences or project instructions** (Claude, Gemini): standing text the AI reads before every chat. Put your role, your business, your preferred tone and format rules there once — every future conversation starts pre-briefed.

- Start your library today with 3 templates from your real life.
- Review monthly: delete what you don't use, refine what you do.
- Ten tested templates quietly outperform any "1000 best prompts" list from the internet — because yours carry your context.

Lesson 3.4 — Creating Beyond Text

The same instruction skills now extend to other media. What advanced users create:

- **Images** — tools like DALL-E (in ChatGPT), Midjourney, and others generate images from text. Describe like a photographer: subject, style, lighting, mood, composition. "Logo sketch for a bakery called Flour & Co — warm minimalist style, flat design, cream and terracotta."

- **Structured documents** — ask for content *as* slides ("outline a 10-slide pitch, with title and 3 bullets per slide"), as spreadsheet formulas, or as a CSV you can import.
- **Code without being a coder** — describe a small tool ("a webpage with a tip calculator") and the AI writes working code, explains how to run it, and fixes errors you paste back. You direct; it types.
- **Voice & video** — text-to-speech, transcription of recordings, and AI video tools all follow the same rule: specific description in, better result out.

The mental shift: you are no longer "chatting." You are **specifying deliverables**. The better your specification (R-T-C-F again), the closer the first draft lands — in any medium.

Lesson 3.5 — Multi-Step Projects

The biggest advanced skill: stop asking AI for a giant result in one prompt, and instead run a **chain of prompts** where each step feeds the next. Big single prompts produce shallow, generic output; chains produce depth.

Example — launching a small online product:

- **Step 1 — Research:** "List the 5 most common customer complaints about existing [product type]. What would 'better' look like?"
- **Step 2 — Strategy:** "Based on those gaps, draft 3 positioning options for my product. Who is each for?"
- **Step 3 — Assets:** "Take positioning #2 and write the landing-page headline, 3 benefit bullets, and an FAQ."
- **Step 4 — Critique:** "Now act as a skeptical customer. Attack this page — every reason you wouldn't buy."
- **Step 5 — Revise:** "Fix the page to answer those objections."

Notice step 4: **making the AI argue against its own work** is one of the most powerful techniques in this entire course. Draft → critique → revise beats a single perfect-prompt attempt every time. Plan the chain before you start: what stages would a professional go through? Then walk the AI through them.

Key Takeaways

- ✓ **R-T-C-F** — Role, Task, Context, Format — is the framework for every prompt that matters.
- ✓ Upload or paste documents and work in stages: summarize → drill down → extract → transform. Say what the output is *for*.
- ✓ Save winning prompts as templates with [placeholders] — your **prompt library** — and set **custom instructions** so every chat starts pre-briefed.
- ✓ Images, slides, spreadsheets, and even working code: you're not chatting, you're **specifying deliverables**.
- ✓ Break big jobs into **prompt chains**, and include a critique step where the AI attacks its own draft.

— PHASE 3 · STANDARD TEST —

Answer all 10 questions, then check yourself against the Answer Key at the back of this book. You need 10/10 to consider the phase passed. If you miss a question, re-read the matching lesson and retake the test — that repetition is exactly how the knowledge sticks.

1. What do the letters in R-T-C-F stand for?

- A) Read, Think, Check, Finish
- B) Role, Task, Context, Format
- C) Request, Time, Cost, Feedback
- D) Role, Topic, Length, Facts

2. What does assigning a **ROLE** ('you are an experienced employment lawyer') actually do?

- A) Legally certifies the AI's answer
- B) Activates the right vocabulary, priorities, and depth for the response
- C) Makes the AI refuse casual questions
- D) Doubles the response length

3. Best practice for a 40-page uploaded report?

- A) One giant prompt asking for everything at once
- B) Ask in stages: overview first, then drill into specific sections
- C) Only read it yourself — AI can't handle documents
- D) Upload it 40 times, one page each

4. Why tell the AI what you'll USE a summary for?

- A) It's required or the AI refuses
- B) The purpose changes what the ideal summary contains
- C) It disables hallucinations
- D) It makes the summary longer

5. What makes a good prompt-library entry?

- A) A screenshot of a good answer
- B) A name, a template with [placeholders], and a note on when to use it
- C) The chat's share link
- D) A list of 1000 prompts copied from the internet

6. What are custom instructions / preferences?

- A) Standing text the AI reads before every chat, so it starts pre-briefed
- B) A way to make the AI respond faster
- C) Rules that block the AI from making mistakes
- D) A paid-only security feature

7. When generating an image, the best prompt style is:

- A) One or two words, to give the AI freedom
- B) Describing like a photographer: subject, style, lighting, mood
- C) Only famous artist names
- D) ALL CAPS for emphasis

8. Why do prompt CHAINS beat one giant prompt for big projects?

- A) They cost less money
- B) Each step's output feeds the next, building depth instead of generic breadth
- C) AI has a strict one-question limit
- D) Chains avoid the need for context

9. In the launch example, what was the purpose of 'act as a skeptical customer and attack this page'?

- A) To test if the AI has emotions
- B) To have the AI expose weaknesses in its own draft so the next revision fixes them
- C) To make the chat more entertaining
- D) To end the conversation politely

10. A non-coder wants a tip-calculator webpage. Realistic AI workflow?

- A) Impossible without a programming degree
- B) Describe the tool, get working code plus run instructions, paste back any errors for fixes
- C) AI can only pseudo-code, never working code
- D) Buy a template — AI can't help with code

Pro

Optimization, troubleshooting & expert frameworks

An advanced user gets great results when things go well. A **pro** gets great results when things go badly — because they can diagnose *why* an output is weak, know the expert techniques to fix it, control hallucination on high-stakes work, pick the right model for each job, and connect AI to their other tools so work happens automatically. Phase 4 is troubleshooting, judgment, and leverage.

Lesson 4.1 — Diagnosing Bad Output

When an answer disappoints, don't just retry — **diagnose**. Almost every weak output has one of four causes, each with its own fix:

- **Vague prompt** → symptoms: generic, could-apply-to-anyone answers. Fix: run R-T-C-F; add specifics and constraints.
- **Missing context** → symptoms: technically correct but wrong *for you*. Fix: supply the situation, the audience, what you've tried, the actual text.
- **Overloaded request** → symptoms: shallow coverage of many things. Fix: split into a chain (Phase 3).
- **Wrong tool or mode** → symptoms: it can't know this (fresh events, your files). Fix: enable web search, upload the document, or switch models.

Two pro moves for stubborn cases: ask the AI **"What information would you need from me to answer this excellently?"** — it will name its own missing context. And ask **"Rate your previous answer's weaknesses, then rewrite it"** — self-critique plus revision, now applied to fixing instead of drafting.

The mindset shift: bad output is a **signal about your input**, not a verdict on AI. Pros read the signal and adjust; beginners retype the same prompt louder.

Lesson 4.2 — Expert Techniques

Five techniques used constantly by expert users, rarely by everyone else:

- **Few-shot examples**: show 2–3 examples of what you want before asking. "Here are two product descriptions in our brand voice: [...] Now write one for this product." Examples beat descriptions of style every time.
- **Think step-by-step**: for logic, math, or tricky decisions, ask the AI to reason through the problem step by step before giving its final answer. Reasoning first, conclusion last,

measurably improves accuracy.

- **Interview me first:** "Before answering, ask me up to 5 questions that would improve your answer." Turns one-shot guessing into consultation.
- **Constraints:** "without jargon", "exactly 3 options", "assume a \$500 budget", "do not suggest hiring anyone." Constraints force useful specificity.
- **Perspective switching:** "Answer as a CFO. Now as a customer. Now as a lawyer." Three lenses reveal what one hides.

None of these require technical skill — they are **management** skills. You're learning to brief, question, and constrain a very fast junior colleague.

Lesson 4.3 — Controlling Hallucination

Phase 1 gave you the Golden Rule (verify important facts). The pro version is a **layered defense** for high-stakes work:

- **Invite honesty:** add "If you're not sure, say so explicitly — do not guess." This measurably reduces confident fabrication.
- **Demand sources:** "cite sources for every factual claim" — then actually open the sources. Fabricated citations look exactly like real ones until clicked.
- **Use web-search modes** for anything recent or factual — grounding answers in live pages beats memory.
- **Cross-examine:** ask the same question in a fresh chat, or to a different AI. Two independent matching answers are far more trustworthy than one.
- **Know the red zones:** exact numbers, dates, quotes, laws, medical dosages, niche facts, and anything about small/local entities are where hallucination concentrates.

Calibrate effort to stakes: brainstorming needs zero verification; a fact going into a contract, a publication, or a business decision needs the full stack. Pros aren't paranoid — they're **calibrated**.

Lesson 4.4 — Choosing Models & Managing Cost

Every provider now offers multiple models, typically a **fast/cheap** tier and a **powerful/reasoning** tier. Choosing well is a pro skill:

- **Fast models** — summaries, rewrites, translations, casual questions, first drafts. Most daily work belongs here.
- **Powerful / reasoning models** — complex analysis, tricky logic, long documents, code, high-stakes writing. Slower and pricier, dramatically better at hard things.
- **Context window** = how much text a model can hold in one conversation. Very long chats degrade — start fresh chats per topic, and paste in only what's relevant.

Cost logic for paid plans: if AI saves you two hours a month, a ~\$20 subscription pays for itself many times over — but only *after* your free-tier usage keeps hitting limits. Upgrade when you feel the ceiling, not before.

One more pro habit: **re-test your assumptions every few months**. Models improve fast; a task AI failed at in January may be trivial by June. Keep a small "failed tasks" note and retry it quarterly.

Lesson 4.5 — Automation: AI That Runs Without You

Until now, every AI interaction started with you typing. The pro frontier is AI that runs **inside your workflows**:

- **Built-in AI in your tools** — Gmail/Outlook drafting replies, Office and Google Docs summarizing and rewriting, Zoom/Teams generating meeting notes. Turn these on first; they're zero effort.
- **Automation platforms** — services like Zapier and Make connect apps with "when X happens, do Y" rules, and AI can be a step: *when a customer email arrives → AI drafts a reply and files it by topic → you review and send*.
- **Scheduled prompts** — some tools can run a prompt on a schedule: a Monday-morning summary of your industry's news, a daily digest of a data source.

The design rule for all automation: **AI drafts, human approves**. Keep a person on every step that sends, spends, deletes, or promises. Automate the preparation, not the responsibility.

Start with exactly one automation this week — the task you dread most that follows a pattern. One working automation teaches you more than ten articles about them.

Key Takeaways

- ✓ Diagnose weak output by cause: vague prompt, missing context, overloaded request, or wrong tool — each has a specific fix.
- ✓ Expert techniques: **few-shot examples, step-by-step reasoning, interview-me-first, constraints, perspective switching**.
- ✓ Control hallucination in layers: invite honesty, demand (and open!) sources, use web search, cross-examine, know the red zones.
- ✓ Match the model to the job — fast for routine, reasoning for hard — and start fresh chats per topic.
- ✓ Automate with the rule **AI drafts, human approves**. Begin with one recurring task this week.

— PHASE 4 · STANDARD TEST —

Answer all 10 questions, then check yourself against the Answer Key at the back of this book. You need 10/10 to consider the phase passed. If you miss a question, re-read the matching lesson and retake the test — that repetition is exactly how the knowledge sticks.

1. An answer is technically correct but wrong for your situation. Most likely cause?

- A) Hallucination
- B) Missing context in your prompt
- C) The model is too small
- D) A billing issue

2. What is 'few-shot' prompting?

- A) Asking several questions at once
- B) Showing 2–3 examples of the desired output before your request
- C) Using as few words as possible
- D) Regenerating a few times and picking the best

3. Why ask the AI to reason step-by-step before answering hard problems?

- A) It makes answers shorter
- B) Working through steps before concluding measurably improves accuracy
- C) It's required for the AI to use math
- D) It prevents all hallucinations

4. The pro way to handle AI-provided sources?

- A) Trust them — citations prove correctness
- B) Open and check them: fabricated citations look real until clicked
- C) Sources are only decoration
- D) Never ask for sources

5. Which task belongs on a FAST model rather than a reasoning model?

- A) Analyzing a 60-page contract for risks
- B) Translating a short casual email
- C) Debugging a subtle logic problem
- D) A high-stakes strategic decision brief

6. What is a context window?

- A) The time limit on a chat
- B) How much text a model can hold in one conversation
- C) The AI's webcam view
- D) The number of users per account

7. The design rule for AI automation is:

- A) Automate everything end-to-end immediately
- B) AI drafts, human approves — keep a person on steps that send, spend, delete, or promise
- C) Never automate anything involving email
- D) Only automate on weekends

8. 'Before answering, ask me up to 5 questions that would improve your answer' is the technique called:

- A) Few-shot prompting
- B) Interview me first
- C) Perspective switching
- D) Chain prompting

9. Where does hallucination risk CONCENTRATE?

- A) Creative brainstorming and fiction
- B) Exact numbers, dates, quotes, laws, and niche/local facts
- C) Grammar corrections
- D) Translations of common phrases

10. Your prompt got a shallow answer covering ten things thinly. Diagnosis and fix?

- A) Vague prompt → add politeness
- B) Overloaded request → split it into a chain of focused steps
- C) Wrong model → always use the biggest one
- D) Missing context → paste more documents

Master

Elite architecture & cutting-edge integrations

Everything so far happened inside a chat window. Mastery is about **architecture**: AI agents that take actions, AI connected to your own knowledge, AI built into products and business processes — and the judgment to deploy all of it responsibly. You won't need to write code in this phase, but you will start thinking like the people who do: in systems, not prompts.

Lesson 5.1 — Agents: AI That Takes Action

A chatbot **answers**; an **agent acts**. Agents are AI systems that can use tools — browse the web, open files, run code, operate applications — and chain many steps toward a goal you set: "research these 5 competitors and produce a comparison report" can now be one instruction, executed over minutes, not a conversation you babysit.

Working with agents well is a management discipline:

- **Brief like a delegator**: goal, constraints, definition of done. "Done means: a table of the 5 competitors' prices, sourced, with links."
- **Set boundaries**: what it may and may not do (spend money, send messages, delete files). Good agent products ask before irreversible steps — insist on that.
- **Review checkpoints**: for long tasks, ask for a plan first, approve it, then let it run. Plan-approve-execute beats fire-and-forget.
- **Verify outputs like a manager**, not a fan: spot-check sources, test the code, scan the edge cases.

The skills transfer perfectly from Phase 4: agents are your expert techniques (constraints, definition of done, review) applied to a worker that moves on its own. People who master delegation to humans adapt fastest to agents — and vice versa.

Lesson 5.2 — Your Own Knowledge Base

By default, AI knows the world but not **your** world. The fix has a technical name — **RAG (Retrieval-Augmented Generation)** — and a simple meaning: connect the AI to your documents, and before answering, it retrieves the relevant pieces and grounds its response on them.

You already can use this without engineering:

- **Projects / spaces** (Claude Projects, custom GPTs, NotebookLM and similar) let you upload a set of documents that every chat in that space can reference: your price list,

policies, past proposals, course notes.

- **A 'company brain'** starts as one folder of clean documents: what you sell, how you talk, your FAQs, your processes. Upload it to a project and every future task starts pre-loaded.
- **Grounded answers cite your documents** — hallucination drops dramatically when the AI quotes your actual policy instead of imagining one.

The master insight: **the quality of your knowledge base becomes the ceiling on the AI's usefulness.** Garbage in, garbage retrieved. Curating clean, current, well-named documents is unglamorous — and it's what separates organizations where AI works from those where it disappoints.

Lesson 5.3 — APIs & Building Without Code

An **API** (Application Programming Interface) is how software talks to software. Every major AI provider sells API access, which means: any app, website, or workflow can send a prompt and receive an answer programmatically. That's how every "AI-powered" product you've seen actually works — a product wrapped around prompts you now know how to write.

You can build on this today, without engineering:

- **No-code app builders** and AI coding assistants let you describe an app in plain language and get a working prototype — the Phase 3 'code without coding' loop, scaled up.
- **Automation platforms** (Phase 4) already speak to AI APIs — your workflows can include an AI step with a prompt template you wrote.
- **System prompts are your R-T-C-F at scale:** when you build anything on an API, you write one master instruction that governs every user interaction. Your prompt-library craft becomes product design.
- **Costs are per-use** (measured in tokens — pieces of words), typically fractions of a cent per request. Prototypes cost pennies.

You don't need to become a developer. You need to know **what's buildable** — because the person who can say "this manual process could be an AI workflow, roughly like this" is the most valuable person in the room, whether or not they write the code.

Lesson 5.4 — Designing AI Workflows for Business

The master framework for putting everything together — turning any real process into an AI-augmented one, in five steps:

- **1. Map the process.** Write every step of the current workflow: who does what, with which inputs and outputs.

- **2. Mark the AI-shaped steps.** Reading, writing, summarizing, classifying, extracting, drafting, translating — anything that transforms information is a candidate. Judgment calls and relationships stay human.
- **3. Insert human checkpoints.** Apply Phase 4's rule at process scale: every irreversible step gets an approver.
- **4. Pilot on real volume.** Run the AI version alongside the old one for a week or a hundred cases. Compare quality, time, and errors honestly.
- **5. Measure and iterate.** Time saved, error rate, satisfaction. Kill what doesn't perform; scale what does.

Worked example — customer support: map (read → categorize → draft → send → log), mark (categorize + draft = AI), checkpoint (human approves sends), pilot (one week, one agent), measure (reply time halved? escalations down?). This template applies identically to invoicing, recruitment screening, content production, and reporting.

Notice what the framework really is: **management thinking with AI literacy** — exactly the two skills this course has been stacking since Phase 1.

Lesson 5.5 — Staying Ahead & Turning Skill Into Value

AI capability roughly doubles on a timescale of months, not years. Mastery is therefore not a state — it's a **practice**:

- **Re-test quarterly.** Your Phase 4 'failed tasks' note is your personal benchmark. When items start passing, capabilities jumped.
- **Evaluate new tools with one question:** what can this do that my current stack can't, on *my* real tasks? Run your own 15-minute trial; ignore hype.
- **Follow releases from the model makers themselves** (OpenAI, Anthropic, Google) — announcements beat commentary.
- **Practice disclosure and ethics:** label AI-generated content where it matters, respect copyright and privacy, and never deploy automation on other people without accountability. Trust is the scarcest resource in the AI economy.

Finally — turning this skill into value. Every organization is currently full of processes waiting for the Phase 5 framework, and short on people who can apply it. Concrete paths: become the AI-workflow person where you work; offer implementation as a freelance service; build a niche AI-assisted product; or teach what you now know to your industry. The window where this skill set is rare is open **right now**.

You came in with no clue. You leave with a system: fundamentals → prompting → frameworks → optimization → architecture. The final test awaits — then the last page of this book gives you your 30-day plan.

Key Takeaways

- ✓ **Agents act** — brief them with goal, constraints, and a definition of done; approve plans before execution; verify like a manager.
- ✓ **RAG / projects** ground AI in *your* documents; your knowledge base's quality is the ceiling on AI's usefulness.
- ✓ **APIs** let any product or workflow call AI; system prompts are your prompt craft becoming product design. Know what's buildable.
- ✓ The 5-step workflow framework: **map** → **mark AI steps** → **human checkpoints** → **pilot** → **measure**.
- ✓ Mastery is a practice: re-test quarterly, trial tools on your real tasks, disclose honestly — and monetize the skill while it's rare.

— PHASE 5 · STANDARD TEST —

Answer all 10 questions, then check yourself against the Answer Key at the back of this book. You need 10/10 to consider the phase passed. If you miss a question, re-read the matching lesson and retake the test — that repetition is exactly how the knowledge sticks.

1. The core difference between a chatbot and an agent?

- A) Agents are always smarter
- B) A chatbot answers; an agent uses tools and takes multi-step actions toward a goal
- C) Agents work only offline
- D) Chatbots are free, agents are not

2. What should a good agent brief include?

- A) Only the topic, to give it freedom
- B) Goal, constraints, and a definition of done
- C) As many tasks as possible at once
- D) A request to skip all confirmations

3. What does RAG do, in plain terms?

- A) Makes the model physically larger
- B) Retrieves relevant pieces of YOUR documents and grounds answers on them
- C) Replaces the need for prompts
- D) Generates random alternative answers

4. Why does document curation matter so much for a knowledge base?

- A) It doesn't — AI fixes messy inputs
- B) The knowledge base's quality becomes the ceiling on the AI's usefulness
- C) Only PDFs can be uploaded
- D) Curation is required by law

5. What is an API?

- A) A type of AI model
- B) A way for software to talk to software — apps sending prompts and receiving answers programmatically
- C) A paid chat interface
- D) An anti-hallucination filter

6. In an API-built product, the system prompt is:

- A) The user's first message
- B) The master instruction governing every interaction — your R-T-C-F craft at product scale
- C) An error log
- D) A password

7. Which steps of a process are the PRIME candidates for AI in the 5-step framework?

- A) Steps that transform information: reading, summarizing, classifying, drafting
- B) Final judgment calls and relationship-building
- C) Steps that spend money irreversibly
- D) All steps, always, immediately

8. Why pilot an AI workflow ALONGSIDE the old process first?

- A) To double everyone's workload permanently
- B) To compare quality, time, and errors honestly before switching
- C) Because AI requires a legal waiting period
- D) To keep the AI secret from the team

9. The best way to evaluate a hyped new AI tool?

- A) Adopt it immediately — hype means quality
- B) Run a 15-minute trial on YOUR real tasks and ask what it does that your stack can't
- C) Wait five years for it to mature
- D) Choose by follower counts

10. Which is a realistic way to monetize this course's skill set?

- A) Selling access to your chat history
- B) Becoming the person who designs AI workflows for a business, as an employee or freelancer
- C) Reselling free-tier accounts
- D) Claiming AI outputs as legal advice

Answer Keys

Grade yourself honestly. Every answer includes a short explanation — read it even when you answered correctly.

PHASE 1 — No Clue

- 1. B** — An LLM generates responses by predicting text from patterns learned in training — it doesn't look answers up in a database, which is exactly why it can hallucinate.
- 2. C** — The chatbot (ChatGPT, Claude...) is the product you interact with; the LLM is the underlying model that powers it — the car and the engine.
- 3. B** — Hallucination = plausible, confidently-stated, false output. It's the #1 weakness to manage, and the reason for the Golden Rule.
- 4. B** — The AI remembers the running thread — a short follow-up in the same chat is exactly how it's meant to be used.
- 5. C** — Models have a training cutoff date. Without a web-search feature turned on, they can't know about fresh events.
- 6. B** — Confidence is not accuracy. Draft freely, verify facts, numbers, names, and quotes before acting on them.
- 7. C** — Public or non-sensitive text is fine. Passwords, card numbers, and other people's personal data never belong in a chat.
- 8. B** — Each conversation is its own context. Some apps add an explicit memory feature, but the default is a fresh start per chat.
- 9. B** — Free tiers are fully usable — paid plans buy better models and more capacity, not correctness guarantees.
- 10. B** — Full sentences, a clear task, a level ('like I'm 15'), and a follow-on request — that's a real instruction, not a search query.

PHASE 2 — Novice

- 1. B** — Task + context + format is the checklist that fixes most bad prompts.
- 2. C** — It names the task, the constraints, the length, and the desired outcome — specific, contextual, formatted.
- 3. B** — Iterating in the same thread with a specific correction is the core novice workflow.
- 4. B** — Being quizzed forces retrieval — the same reason this book ends each phase with a test.

- 5. C** — Photo + question ('what is this and what should I do?') is one of the most useful mobile skills.
- 6. B** — A couple of refinement rounds is the workflow working as intended, not a failure.
- 7. A** — If you don't state a format, the AI guesses. Asking for a table with named columns gets you exactly the layout you want.
- 8. A** — Voice mode shines hands-free: language practice, rehearsing conversations, spoken explanations. It doesn't change accuracy or security.
- 9. B** — Setting the level tells the AI how much prior knowledge to assume — a core tutoring technique.
- 10. B** — Context shapes tone, formality, and content. The AI can only tailor what it knows about.

PHASE 3 — Advanced

- 1. B** — Role, Task, Context, Format — supply all four for any prompt where quality matters.
- 2. B** — A role steers the perspective and expertise the model draws on — it doesn't certify or guarantee anything.
- 3. B** — Staged questioning (overview → drill-down) reliably beats one enormous request.
- 4. B** — A summary for a 30-second CEO brief and one for deciding conference attendance select different information.
- 5. B** — Name + reusable template + usage note makes prompts findable and reusable. Your 10 tested templates beat generic lists.
- 6. A** — They're your permanent briefing: who you are, your tone and format rules — applied automatically to every conversation.
- 7. B** — Specific visual description (subject, style, lighting, composition) is the image equivalent of R-T-C-F.
- 8. B** — Staged work mirrors how professionals think: research → strategy → draft → critique → revise. Depth compounds.
- 9. B** — Draft → self-critique → revise is one of the most powerful patterns in the course.
- 10. B** — Describe → generate → run → paste errors back → fix. You direct; the AI types.

PHASE 4 — Pro

- 1. B** — Correct-but-not-for-you is the signature of missing context: the AI answered the generic version of your question.
- 2. B** — Examples communicate style and format better than any description of them.

- 3. B** — Reasoning first, conclusion last gives the model room to work — accuracy improves on logic-heavy tasks.
- 4. B** — Hallucinated citations are formatted perfectly. A source only counts once you've opened it.
- 5. B** — Routine rewrites, translations, and summaries are fast-model work; save reasoning models for genuinely hard tasks.
- 6. B** — It's the model's working memory for the conversation — why very long chats degrade and fresh chats help.
- 7. B** — Automate preparation, not responsibility. Human review guards every irreversible step.
- 8. B** — It converts one-shot guessing into consultation — the AI names its own missing context.
- 9. B** — Precise, verifiable, niche facts are the red zone — calibrate verification effort there.
- 10. B** — Shallow-everywhere is the overload signature. Chains (Phase 3) restore depth step by step.

PHASE 5 — Master

- 1. B** — Action + tool use + multi-step autonomy is what makes it an agent.
- 2. B** — Brief like a delegator. 'Done means...' is the master's phrase.
- 3. B** — Retrieval-Augmented Generation = fetch from your knowledge base first, then answer from it — which also slashes hallucination.
- 4. B** — Garbage in, garbage retrieved. Clean, current, well-named documents are the unglamorous foundation.
- 5. B** — API access is how every 'AI-powered' product works under the hood.
- 6. B** — One standing instruction shapes all user interactions — prompt writing becomes product design.
- 7. A** — Information-transforming steps are AI-shaped; judgment and relationships stay human; irreversible steps get human checkpoints.
- 8. B** — A side-by-side pilot on real volume is the honest test — measure, then scale or kill.
- 9. B** — Your tasks are the only benchmark that matters. Announcements from model makers beat commentary; trials beat both.
- 10. B** — Process mapping + AI literacy is exactly the scarce combination organizations are paying for right now.

You Made It — Now What?

If you completed all five tests honestly, you are now more capable with AI than the vast majority of people — including many who work in tech. But mastery is a moving target: the tools will keep changing, and the skills in this book are what let you keep up.

Your 30-day plan from here: (1) Pick one recurring task in your life or work and fully delegate it to AI this week. (2) Build your prompt library to 10 tested templates. (3) Automate one workflow end-to-end. (4) Teach someone else the R-T-C-F framework — teaching is the fastest way to lock in mastery.

The people who win with AI are not the ones who know the most models. They are the ones who built the habit of asking: "**Could AI do this step for me?**" — and now, that's you.

Thank you for learning with this course. Now go build something.

© 2026 Paper Plane Workshop. Free edition — you are welcome to share the download link (paperplaneworkshop.com), but this book may not be resold or modified. The information in this book is provided for educational purposes; AI tools change quickly, and specific product features may differ by the time you read this.