

How to Use Claude

A Complete Beginner's Practical Tutorial

Verified 21 August 2026

Estimated completion time: 3-4 hours, including exercises and the practical project

Difficulty level: Complete beginner

Tutorial scope verified: 21 August 2026

Before you begin

1. Assumed learner profile

This tutorial assumes that the learner:

- Has little or no previous knowledge of artificial intelligence.
- Has never used Claude.
- Does not know programming, mathematics, or data science.
- Can use a web browser, enter text, open files, and follow menu instructions.
- Wants to use Claude for everyday learning, writing, research, planning, document analysis, and professional tasks.

2. Scope of this tutorial

This tutorial covers the **Claude consumer application**, primarily through its browser interface. It explains how to:

- Create and configure an account.
- Hold useful conversations with Claude.
- Write effective prompts.
- Improve answers through follow-up instructions.
- Upload and analyse documents and images.
- Use web search for current information.
- Organise work using Projects.
- Create simple Artifacts.
- Check Claude's answers before using them.
- Protect privacy and use AI responsibly.

3. Intentionally excluded

The following advanced subjects are outside this beginner tutorial:

- The Claude API.
- Claude Code and software-development workflows.
- Programming applications using Claude.
- Model training or fine-tuning.
- Mathematical explanations of neural networks.
- Model Context Protocol, or MCP.
- Enterprise administration and organisation-wide deployment.
- Autonomous agents and automated actions.
- Detailed comparisons with ChatGPT, Gemini, Copilot, or other AI products.

These should be studied only after the learner is comfortable with the basic Claude application.

4. Tools and versions used

Tool	Version or assumption
Claude	Current web application as of 21 August 2026
Default model	Claude Sonnet 5 where available
Optional advanced model	Claude Opus 5, depending on the learner's plan
Browser	A current version of Chrome, Edge, Safari, or Firefox
Device	Desktop or laptop recommended; mobile also supported
Account	Free Claude account is sufficient for the main tutorial

Tool	Version or assumption
Programming software	Not required

Claude Sonnet 5 is the default model on the Free and Pro plans, while Claude Opus 5 is available on selected paid plans. Model availability can vary by subscription, organisation settings, and product rollout. [Anthropic: Claude Sonnet 5](#)

Learning outcomes

By the end of this tutorial, the learner should be able to:

1. Explain Claude, generative AI, and large language models in simple terms.
2. Distinguish between the Claude application, a Claude model, a prompt, a chat, a Project, and an Artifact.
3. Create an account and configure basic privacy settings.
4. Write clear prompts using a repeatable structure.
5. Improve weak or incomplete responses through follow-up instructions.
6. Upload documents and ask evidence-based questions about them.
7. Use web search appropriately for current information.
8. Create a basic Claude Project.
9. Complete a realistic meeting-notes-to-action-plan project.
10. Identify hallucinations, privacy risks, security concerns, bias, and overreliance.
11. Validate important outputs before using or sharing them.

Prerequisites

Required equipment and access

You need:

- A computer, tablet, or smartphone.
- A stable internet connection.
- A modern web browser.
- An email address.
- Access to a supported country or region.
- A supported mobile number if phone verification is requested.

Claude is available through its website, Windows and macOS desktop applications, and iOS and Android applications. Anthropic states that users must be at least 18 years old and located in a supported region.

[Anthropic: Getting started with Claude](#)

Software installation

No installation is required for this tutorial. The browser version is sufficient.

To use the optional desktop application:

1. Open Claude's official download page.
2. Select **Windows** or **macOS**.
3. Download the installer.
4. Open the downloaded file.
5. Follow the installation instructions.
6. Launch Claude and sign in.

For mobile access:

- On iPhone or iPad, search for **Claude by Anthropic** in the Apple App Store.

-
- On Android, search for **Claude by Anthropic** in Google Play.
 - Confirm that the publisher is Anthropic before installing. [Anthropic: Install Claude Desktop](#)

Subscription

The free plan is adequate for:

- Basic chats.
- Prompting practice.
- Document analysis within plan limits.
- Web search where available.
- A limited number of Projects.
- Basic Artifacts.

Paid plans generally provide higher usage allowances and additional model options. Do not purchase a paid plan merely to complete this tutorial.

Table of contents

1. [What Claude is and why it matters](#)
2. [Core concepts and terminology](#)
3. [Creating an account and configuring privacy](#)
4. [Understanding the Claude interface](#)
5. [Starting your first useful conversation](#)
6. [Writing effective prompts](#)
7. [Improving answers through follow-up prompts](#)
8. [Uploading and analysing files and images](#)
9. [Using web search, Research, and extended thinking](#)
10. [Using Projects, memory, and Artifacts](#)
11. [Practical project: Meeting notes to verified action plan](#)
12. [Visual and screenshot specifications](#)
13. [Diagrams and graph decisions](#)
14. [Exercises](#)
15. [Solutions and model answers](#)
16. [Troubleshooting guide](#)
17. [Responsible AI considerations](#)
18. [Summary checklist](#)
19. [Recommended next topics](#)
20. [Glossary](#)
21. [References](#)

1. What Claude is and why it matters

1.1 A simple explanation

Claude is an artificial intelligence assistant developed by Anthropic.

You communicate with it by entering instructions or questions in ordinary language. Claude then generates a response based on:

- Your instruction.
- Earlier messages in the current conversation.
- Files or images you provide.
- Information obtained through enabled tools such as web search.
- Patterns learned during its training.

You can ask Claude to:

- Explain a difficult subject.
- Summarise a document.
- Rewrite an email.
- Generate ideas.
- Compare options.
- Analyse an image.
- Organise meeting notes.
- Develop a learning plan.
- Review an argument.
- Draft tables, reports, or presentations.
- Help write or explain code.

Anthropic describes Claude as a large language model designed to act as a helpful conversational assistant. Its documented use cases include writing, learning, summarisation, brainstorming, translation, coding, and image analysis. [Anthropic: Getting started with Claude](#)

1.2 What “artificial intelligence” means

Artificial intelligence, usually shortened to **AI**, is a broad term for computer systems that perform tasks commonly associated with human intelligence.

Examples include:

- Recognising speech.
- Identifying objects in an image.
- Translating languages.
- Recommending products.
- Generating text.
- Finding patterns in data.

AI does not necessarily think or understand in the same way a person does. Different AI systems perform different tasks using different methods.

1.3 What “generative AI” means

Generative AI is AI that creates new content.

It may generate:

- Text.
- Images.
- Audio.
- Video.

- Software code.
- Diagrams.
- Structured data.

Claude is primarily a text-based generative AI assistant, although it can also analyse images and create certain visual or interactive outputs.

1.4 What a large language model is

A **large language model**, or **LLM**, is an AI model trained on very large collections of text and other data.

Simplified explanation

You can imagine an LLM as an extremely advanced text-completion system. It examines your input and calculates what response is likely to be relevant and useful.

More precise explanation

An LLM is a neural-network-based statistical model that estimates relationships among units of language and generates sequences in response to instructions and context.

It does not normally retrieve a stored answer in the way a traditional database does. It generates an answer from learned patterns unless it is using an additional tool, such as web search or file retrieval.

1.5 Why Claude matters

Claude can reduce the time required to produce a first draft, explore a topic, structure information, or examine alternatives.

For example:

Situation	Possible use of Claude
You receive a long report	Ask for a structured summary with page references
You need to write an email	Provide the facts and ask for a professional draft
You are learning a new topic	Ask for a beginner explanation followed by a quiz
You have meeting notes	Extract decisions, owners, deadlines, and open questions
You are comparing options	Request a criteria-based comparison table
You need current information	Enable web search and request authoritative sources
You are developing an argument	Ask Claude to identify assumptions and counterarguments
You need a plan	Ask for stages, dependencies, risks, and success measures

1.6 What Claude does not guarantee

Claude does not guarantee that its answers are:

- Factually correct.
- Current.
- Complete.
- Unbiased.
- Legally compliant.
- Professionally appropriate.
- Supported by real sources.
- Safe to use without review.

Claude can produce convincing but incorrect statements. Anthropic refers to these errors as hallucinations and advises users to verify important information against reliable sources. [Anthropic: Incorrect or misleading responses](#)

Knowledge check 1

1. What is the difference between AI and generative AI?
2. Is Claude a database that always retrieves a stored answer?
3. Name three useful tasks Claude can help with.
4. Why should an important Claude response be checked?

2. Core concepts and terminology

2.1 The Claude application and the Claude model

These terms are related but not identical.

Claude application

The **application** is the interface you open in a browser, desktop program, or mobile app.

It contains features such as:

- Chat history.
- File attachment.
- Web search.
- Projects.
- Artifacts.
- Settings.
- Model selection.

Claude model

The **model** is the AI system that processes your prompt and generates a response.

A simple analogy:

- The application is the office.
- The model is the specialist working inside the office.
- Your prompt is the work brief.
- The answer is the delivered work.

2.2 Prompt

A **prompt** is the instruction, question, information, or material you give Claude.

Example:

Explain ransomware in five simple bullet points for someone who has never studied cybersecurity.

A prompt can contain:

- A task.
- Background information.
- Source text.
- Rules.
- Examples.
- A required output format.

2.3 Response or output

A **response**, also called an **output**, is what Claude generates after receiving your prompt.

The quality of the response depends partly on:

- How clearly you explained the task.
- Whether you supplied enough relevant context.
- Whether your requirements conflict.
- Whether Claude has access to reliable source material.
- Whether the task requires current information.
- Whether you verify and refine the answer.

2.4 Chat or conversation

A **chat** is an ongoing sequence of prompts and responses.

Claude normally uses earlier messages in the same conversation as context. This lets you say:

Make the second paragraph shorter.

Claude can infer that you are referring to the second paragraph of the previous response.

However, very long conversations may become inefficient or confused. Starting a new chat with a concise summary can sometimes produce a better result.

2.5 Context

Context is the information Claude can consider while generating a response.

It may include:

- Your current prompt.
- Earlier messages.
- Uploaded documents.
- Project instructions.
- Project knowledge.
- Information obtained through enabled tools.

A useful analogy is a work desk:

- The material on the desk is the current context.
- Material stored somewhere else may not be visible.
- When too much material accumulates, important details can become harder to locate.

2.6 Token

A **token** is a small unit used by a language model to process text.

A token may be:

- A whole short word.
- Part of a longer word.
- Punctuation.
- A number.
- A symbol.

You do not need to count tokens manually as a beginner. The important point is that long conversations, lengthy files, complex tasks, and extensive outputs consume more of the available usage and context capacity. Claude's usage and context limits vary by plan, selected model, feature, and conversation complexity. [Anthropic: Usage and length limits](#)

2.7 Knowledge cutoff

A **knowledge cutoff** is the approximate latest date included in a model's original training knowledge.

A model may not know events that happened after its cutoff unless it uses a current external source such as web search.

Anthropic currently lists different reliable knowledge cutoffs for different Claude models. For example, Claude Opus 5 has a later listed cutoff than Claude Sonnet 5. A newer model can still lack information about events that occurred after its training data was prepared. [Anthropic: Training data recency](#)

Therefore:

- Use ordinary chat for general explanations and writing.
- Use web search for recent events, prices, laws, product details, schedules, or changing information.
- Check the cited source yourself before making an important decision.

2.8 Hallucination

A **hallucination** occurs when an AI generates false or unsupported information that sounds plausible.

Examples include:

- A nonexistent academic paper.
- An incorrect law.
- A fabricated quotation.
- A wrong calculation.
- An invented feature.
- An inaccurate historical date.
- A summary that adds information not found in the source.

A hallucination is not necessarily a deliberate lie. It results from how a generative model constructs language.

2.9 Project

A **Project** is a dedicated workspace for related chats, documents, and instructions.

You might create separate Projects for:

- A university course.
- A business proposal.
- A job search.
- A research topic.
- A personal learning programme.

Project knowledge and instructions help maintain consistency across chats within the same workspace.

2.10 Artifact

An **Artifact** is content Claude creates in a separate working panel.

Examples include:

- A formatted document.
- A diagram.
- A simple interactive tool.
- An HTML page.
- A visual planner.
- A small application.
- A structured report.

An Artifact is different from a normal chat answer because it can be viewed and refined as a separate object.

2.11 Tool

A **tool** gives the model an additional capability.

Examples include:

- Web search.

- File analysis.
- Code execution.
- Integrations with external services.

The model should not be assumed to have used a tool merely because it claims that it did. Look for visible evidence, such as citations, attached-file references, or an actual completed action.

Diagram: Basic Claude workflow




Figure 1. Basic Claude workflow.

The process should not end immediately after Claude generates an answer. Human review and revision are part of responsible AI use.

Knowledge check 2

1. What is the difference between the Claude application and a Claude model?
2. What is context?
3. What is a hallucination?
4. When should you use web search rather than relying only on the model?
5. What is the difference between a Project and a normal chat?

3. Creating an account and configuring privacy

Stage purpose

This stage gives you access to Claude and helps you make basic privacy decisions before entering personal or work information.

3.1 Open Claude

1. Open your web browser.
2. Navigate to Claude's official website.
3. Check that the domain is the official Claude domain before entering account information.
4. Select the option to create an account or continue to the sign-in screen.

Expected result: You should see an account screen offering supported sign-in methods.

Visual placement: Insert Visual 2, "Actual Claude sign-in page," here.

3.2 Create your account

Depending on the current sign-up screen:

1. Enter your email address and select **Continue with email**, or select **Continue with Google**.
2. Follow the instructions shown on screen.
3. When using email, open the message sent by Claude.
4. Select the secure sign-in or verification link.
5. Complete any requested profile or terms steps.

Anthropic's account guidance currently documents email and Google sign-in flows. [Anthropic: Sign up for Claude](#)

Phone verification

A new account may require mobile-number verification.

1. Select the correct country code.
2. Enter your mobile number.
3. Request the verification code.
4. Enter the code received by SMS.
5. Complete the remaining account steps.

If the code does not arrive:

- Recheck the number and country code.
- Wait at least five minutes.
- Select **Try again**.
- Use a standard mobile number supported by the service.
- Do not repeatedly request codes in rapid succession.

Anthropic states that supported phone verification may be required for new accounts and that verification cannot simply be skipped when requested. [Anthropic: Verify your phone number](#)

Expected result: Claude should open its main home or chat screen.

3.3 Review model-improvement settings

Before entering personal material:

1. Select your profile initials or account menu.
2. Open **Settings**.
3. Open **Privacy**.
4. Locate **Help Improve our AI Models**, or the nearest equivalent label.

5. Read the description.
6. Decide whether to allow your consumer conversations to be used for model improvement.
7. Change the toggle according to your privacy preference.

The exact wording or location may change during interface updates. Anthropic provides consumer users with controls over the use of conversations for model improvement. [Anthropic: Model-improvement privacy settings](#)

Recommended beginner practice

For personal experimentation:

- Review the setting rather than accepting it without reading.
- Avoid entering passwords, bank information, health records, private identity documents, or confidential corporate information.

For professional or organisational work:

- Follow your employer's approved-AI policy.
- Use the organisation's managed account where required.
- Do not assume that a personal consumer account is authorised for corporate data.
- Classify and sanitise information before uploading it.

Visual placement: Insert Visual 7, "Claude privacy settings," here.

3.4 Understand Incognito chats

An **Incognito chat** is designed for a temporary conversation that does not appear in normal chat history or memory and is not used for model training.

To start one, where available:

1. Start a new chat.
2. Select the ghost-shaped Incognito control near the upper-right area.
3. Confirm that the interface indicates Incognito mode.
4. Enter your prompt.

Important limitations:

- Incognito does not mean that no data is processed or retained.
- Anthropic states that Incognito chats are normally retained for 30 days.
- They cannot be converted into regular saved chats.
- Organisational retention and export rules may differ for managed plans. [Anthropic: Use Incognito chats](#)

Use Incognito as an additional privacy control, not as permission to submit secrets.

3.5 Understand deletion

Deleting a consumer chat removes it from your visible history immediately. Anthropic states that deleted chat data is generally removed from backend systems within 30 days, subject to specified safety, legal, or retention exceptions. Data previously selected for model improvement may be governed by longer de-identified retention conditions. [Anthropic: Data retention](#)

The practical lesson is simple:

Do not upload sensitive data on the assumption that deleting the conversation later will erase every risk.

Knowledge check 3

1. Does Incognito mean that a conversation is never retained?
2. Should you upload confidential company information through a personal account merely because you turned off model improvement?
3. Where should you look for the model-improvement control?
4. What should you do if an SMS code has not arrived after several minutes?

4. Understanding the Claude interface

Stage purpose

This stage introduces the main controls required for everyday use.

The exact positions and labels can change as Anthropic updates the product. Focus on the function of each control rather than memorising its location.

Visual placement: Insert Visual 3, “Claude main chat interface,” here.

4.1 Main interface elements

Sidebar

The sidebar usually provides access to:

- New chat.
- Previous chats.
- Projects.
- Artifacts.
- Search or history.
- Account settings.

Prompt box

The prompt box is where you type your instruction.

It may also contain controls for:

- File attachment.
- Web search.
- Model selection.
- Effort or thinking level.
- Sending the message.

Chat area

The chat area displays:

- Your prompts.
- Claude’s responses.
- Citations.
- File references.
- Tool activity.
- Artifacts or generated files.

Model menu

The current model and effort controls may appear beside the send button.

To inspect the available models:

1. Open a new chat.
2. Locate the current model name near the prompt box or send button.
3. Select it.
4. Review the available models.
5. Keep the default model for ordinary beginner tasks.

Anthropic’s current interface allows eligible users to select a model, effort level, and, for suitable tasks, extended thinking. Higher effort may use more time and usage capacity. [Anthropic: Model, effort, and thinking settings](#)

4.2 Which model should a beginner choose?

Use the default model for:

- General questions.
- Email writing.
- Summaries.
- Brainstorming.
- Document extraction.
- Everyday explanations.

Use a stronger model, where available, for:

- Complex analysis.
- Detailed criticism.
- Difficult reasoning.
- Large document synthesis.
- Complicated planning.

A stronger model does not remove the need for verification.

4.3 Effort or thinking settings

Some accounts show an effort setting such as:

- Low.
- Medium.
- High.
- Maximum or equivalent.

Use:

- **Default or medium effort** for most work.
- **Higher effort** for difficult analysis, planning, or reasoning.
- **Lower effort** for quick reformatting or simple rewriting.

Higher effort does not automatically improve a poor prompt. Clear instructions and reliable evidence remain essential.

Knowledge check 4

1. Where do you type an instruction to Claude?
2. Should a beginner always choose the most powerful model and highest effort?
3. What is the purpose of the sidebar?
4. Does selecting a stronger model guarantee a correct answer?

5. Starting your first useful conversation

Stage purpose

This stage shows how to ask a simple question, inspect the result, and use a follow-up instruction.

5.1 Start a new chat

1. Select **New chat**.
2. Click inside the prompt box.
3. Enter the following prompt:

Explain two-factor authentication in five simple bullet points.

Assume I have no technical knowledge.

Finish with one everyday example.

4. Select the send button or press the appropriate keyboard shortcut.

Expected output

Claude should produce:

- Five brief explanatory bullet points.
- Plain language.
- An example such as using a password plus a phone code.

The exact wording will vary because generative AI outputs are not always identical.

5.2 Ask a follow-up question

In the same conversation, enter:

Now explain the difference between a password and the second factor using a house-security analogy.

Expected output

Claude should use the earlier subject as context and produce an analogy, for example:

- Password = the house key.
- Second factor = an additional alarm code or identity check.

This demonstrates that you do not need to repeat the entire subject in every follow-up message.

5.3 Correct the level of difficulty

If the answer is still too technical, enter:

This is still too technical for a complete beginner.

Rewrite it using shorter sentences, define every technical term, and do not use abbreviations without explaining them.

Expected output

Claude should revise the explanation rather than merely repeat it.

5.4 Change the format

Enter:

Convert the explanation into a two-column table:

Column 1: Concept
Column 2: Simple explanation

Claude should present the same underlying information in a table.

5.5 End with a self-test

Enter:

Give me three multiple-choice questions to test whether I understood the topic.

Do not show the answers until after I respond.

This turns a single explanation into an interactive learning activity.

Troubleshooting

Problem	Correction
The answer is too long	Specify a word or bullet limit
The answer is too technical	State the learner level and request definitions

Problem	Correction
The answer is too general	Add context and a concrete purpose
Claude changes the topic	Restate the objective in one sentence
The format is wrong	Provide an exact heading, table, or template
The response stops early	Ask Claude to continue from the final completed section
The answer seems false	Request sources and independently verify them

Knowledge check 5

1. Why can you refer to “the explanation” in a follow-up message?
2. How can you make an answer less technical?
3. What instruction can control the response length?
4. Why might the exact output differ when the same prompt is repeated?

6. Writing effective prompts

Stage purpose

This stage gives you a repeatable method for writing prompts that are clearer, more useful, and easier to evaluate.

Anthropic recommends giving Claude clear, direct instructions, relevant context, explicit constraints, and a defined output format. It also recommends breaking sequential work into ordered steps and providing examples when consistency matters. [Anthropic: Prompt library](#)

6.1 The G-C-C-F-C prompt structure

The following is an instructional framework used in this tutorial. It is not an official Anthropic acronym.

G: Goal

What should Claude do?

Examples:

- Explain a concept.
- Summarise a document.
- Draft an email.
- Compare two options.
- Identify risks.
- Produce a plan.

C: Context

What background does Claude need?

Examples:

- Who the audience is.
- Why the task matters.
- What has already happened.
- What source material should be used.
- What role you have.
- What decision the output will support.

C: Constraints

What limits or rules must Claude follow?

Examples:

- Maximum 300 words.
- Use only the attached file.
- Do not invent missing information.
- Use plain English.
- Include no confidential details.
- Distinguish facts from assumptions.
- Use official sources only.

F: Format

How should the result be presented?

Examples:

- Table.
- Email.
- Checklist.
- Numbered plan.
- Executive summary.
- Frequently asked questions.
- JSON or code, for technical users.

C: Check

How should Claude inspect its own work?

Examples:

- List assumptions.
- Flag uncertainty.
- Check calculations.
- Cite supporting evidence.
- Identify missing information.
- State what requires human verification.

6.2 Prompt anatomy diagram




Figure 2. Five-part prompt anatomy.

Visual placement: Insert Visual 4, “Five-part prompt anatomy,” here.

6.3 Weak prompt versus stronger prompt

Weak prompt

Write something about cybersecurity.

Problems:

- The purpose is unclear.
- The audience is unknown.
- The length is unspecified.
- The required format is unknown.

- “Cybersecurity” is too broad.
- There is no quality check.

Stronger prompt

```
Goal:
Write a beginner-friendly explanation of phishing.

Context:
The audience is office employees with no technical background.
The explanation will be used in a five-minute awareness briefing.

Constraints:
- Maximum 250 words.
- Use plain English.
- Define phishing the first time it appears.
- Include one email example.
- Do not use fear-based language.
- Do not claim that every suspicious email is malicious.

Format:
1. Definition
2. Three warning signs
3. One example
4. What the employee should do

Quality check:
At the end, identify any advice that may depend on the organisation's internal security policy.
```

This prompt gives Claude a usable work specification.

6.4 You do not always need labels

The labels are useful while learning, but a natural paragraph can work equally well.

Example:

```
Rewrite the email below for a supplier. Keep the tone firm but professional,
use no more than 180 words, preserve all dates and amounts, and do not add
commitments that are not in the original. Provide a subject line followed by
the email body. After the draft, list any unclear points that should be
confirmed before sending.
```

The structure still contains:

- Goal.
- Context.
- Constraints.
- Format.
- Check.

6.5 Give Claude the source of truth

When accuracy matters, state which material controls the answer.

Example:

```
Use the attached proposal as the only source of truth.

Do not use general knowledge to fill gaps.

For every commercial commitment you identify, provide:
- the exact commitment;
- the relevant page number;
- whether it is explicit or inferred;
- any ambiguity requiring clarification.
```

This reduces, but does not eliminate, the risk of invention.

6.6 Tell Claude what not to do

Negative constraints are useful when they target a specific risk.

Examples:

```
Do not invent missing owners or deadlines.
Do not present assumptions as confirmed facts.
Do not cite a source unless you have actually used it.
Do not change the meaning of the original text.
```

Avoid a very long list of prohibitions. Too many rules can conflict and reduce clarity.

6.7 Ask for an example

Examples are particularly useful when you want consistent style or format.

```
Use the following format for every risk:

Risk:
Cause:
Potential impact:
Existing control:
Recommended action:
Owner:
Target date:

Example:
Risk: Delayed venue confirmation
Cause: Booking has not been approved
Potential impact: Marketing cannot begin
Existing control: Weekly planning review
Recommended action: Escalate the booking request
Owner: Events coordinator
Target date: 2026-08-25
```

An example shows Claude what an acceptable output looks like.

6.8 Use roles carefully

A prompt may say:

```
Act as an experienced project manager.
```

This can help establish perspective and terminology. It does **not**:

- Give Claude a real professional licence.
- Make its claims authoritative.
- Replace an actual project manager.
- Guarantee that it follows a particular standard.
- Give it access to facts you did not provide.

A better role prompt includes scope and evidence requirements:

```
Review this plan from the perspective of an experienced project manager.

Focus on dependencies, ownership, schedule risk, acceptance criteria, and
unresolved decisions.

Do not claim that the plan complies with a named standard unless the supplied
evidence demonstrates compliance.
```

6.9 Ask for alternative views

For important decisions, do not ask Claude only to support your preferred answer.

Use:

```
Present the strongest case for this proposal.

Then present the strongest credible case against it.

Identify the assumptions on which each case depends.

Recommend the conditions under which the proposal should proceed, change,
or stop.
```

This reduces confirmation bias.

Knowledge check 6

1. What does each letter in G-C-C-F-C represent?
2. Why should you identify the source of truth?
3. Does telling Claude to “act as a lawyer” make it a licensed lawyer?
4. When is an example useful in a prompt?
5. Why should important decisions include counterarguments?

7. Improving answers through follow-up prompts

Stage purpose

Good Claude use is usually iterative. The first answer should often be treated as a draft rather than a final product.

7.1 The four-pass method

Pass 1: Generate

Ask for the initial result.

```
Create a project plan from the information below.
```

Pass 2: Audit

Ask Claude to inspect weaknesses.

```
Audit your plan.
```

```
Identify:
```

- unsupported assumptions;
- missing owners;
- missing deadlines;
- unclear dependencies;
- conflicting requirements;
- risks not supported by the source material.

```
Do not revise it yet.
```

Pass 3: Correct

Give Claude your corrections.

```
Apply the following corrections:
```

- The project owner is Lina.
- The budget is not yet approved.
- The target date is tentative.
- Remove any statement that suggests procurement has started.

Pass 4: Finalise

Request a clean output.

```
Produce the final plan without the discussion history.
```

```
Use:
```

1. Objective
2. Scope
3. Milestones
4. Responsibilities
5. Dependencies
6. Risks
7. Open decisions

7.2 Ask Claude to expose uncertainty

Useful instructions include:

```
Separate confirmed facts, reasonable inferences, assumptions, and recommendations.  
For each conclusion, state the supporting evidence and your confidence.  
Identify any information you could not confirm.  
Where the evidence is insufficient, write "Not established" rather than guessing.
```

Claude's confidence statement is not a mathematical guarantee. It is a communication aid that still requires human checking.

7.3 Ask for traceability

Traceability means being able to connect an output back to its source.

Example:

```
Create a table with these columns:  
1. Finding  
2. Supporting text  
3. File name  
4. Page or section  
5. Fact, inference, or assumption  
6. Verification required
```

This is more useful than requesting a summary with no evidence trail.

7.4 Correct Claude directly

When Claude makes a mistake, do not merely write:

```
Wrong.
```

Use:

```
The deadline in your table is incorrect.  
The source states 25 August, not 28 August.  
Correct the table, then check every other date against the source before  
producing the revised version.
```

Specific feedback teaches the model what to change inside that conversation.

7.5 Start a new chat when necessary

Start a new chat when:

- The subject has changed substantially.
- The conversation contains many abandoned instructions.
- Claude repeatedly follows an old requirement.
- The chat has become very long.
- You need a clean, reproducible result.
- You are moving from brainstorming to controlled final production.

Before starting a new chat, ask:

```
Create a concise handover summary for a new conversation.  
  
Include:  
- objective;  
- confirmed facts;  
- decisions;  
- unresolved questions;  
- final constraints;  
- required output;  
- information that must not be assumed.
```

Review the handover yourself before pasting it into the new chat.

7.6 Use concise rather than repetitive correction

Avoid continuously appending conflicting instructions such as:

- Be detailed.
- Make it short.
- Add more evidence.
- Remove the details.
- Keep everything.
- Use only three bullets.

Instead, replace them with one consolidated instruction:

Produce a maximum 500-word executive brief.

Include only the three highest-impact findings, but provide one supporting piece of evidence and one action for each.

Knowledge check 7

1. What are the four passes in the four-pass method?
2. Why is "Wrong" weaker feedback than a specific correction?
3. What does traceability mean?
4. When might starting a new chat improve the result?
5. Is Claude's stated confidence a guarantee of correctness?

8. Uploading and analysing files and images

Stage purpose

This stage shows how to give Claude source material and ask questions that remain tied to that material.

Claude supports file uploads in chats and Projects. Current documented chat limits include up to 20 files in a conversation, up to 500 MB per file, images up to 8,000 x 8,000 pixels, and PDFs up to 1,000 pages, although actual processing also depends on plan and available context. [Anthropic: Supported document uploads](#)

8.1 Prepare a safe file

Before uploading:

1. Confirm that you have permission to use the file.
2. Remove passwords, account numbers, identity documents, personal records, and unnecessary confidential data.
3. Check whether your employer permits the use of the selected Claude account.
4. Use a copy rather than the only original.
5. Give the file a clear name.

Example:

Workshop_Budget_Draft_v1.pdf

Avoid:

document-final-new-final2.pdf

8.2 Attach the file

1. Open a new chat.
2. Select the attachment or plus control beside the prompt box.
3. Choose the file.

4. Wait until the file appears as attached.
5. Do not send the prompt until the upload is complete.

Visual placement: Insert Visual 6, "Uploading a document to Claude," here.

Expected result: The file name should appear in the prompt area or attachment list.

8.3 Confirm that Claude can identify the file

Enter:

```
List the files attached to this conversation.

For each file, state:
- file name;
- apparent document type;
- number of pages or sections you can identify;
- any readability problems.

Do not analyse the content yet.
```

Expected output

Claude should identify the uploaded file and warn about obvious problems such as:

- Unreadable scans.
- Missing pages.
- Image-only text.
- Complex tables.
- Unsupported or damaged content.

Do not assume that a successful upload means every part of the file was correctly understood.

8.4 Request an evidence-based summary

Use:

```
Use only the attached document as the source of truth.

Create a summary with these sections:
1. Purpose of the document
2. Main commitments
3. Dates and deadlines
4. Costs or financial terms
5. Responsibilities
6. Risks and ambiguities

For every material statement, provide the page or section where it appears.

Do not fill missing information using general knowledge.
Mark unclear information as "Unclear in source."
```

Expected output

You should receive a structured summary with document locations.

Check at least three cited pages yourself.

8.5 Analyse a table

For a financial or operational table:

```
Extract the table named "Annual Costs."

Preserve:
- row labels;
- units;
- currencies;
- decimal places;
- totals;
- footnotes.
```

```
Then independently recalculate each subtotal and the final total.
Show the formula used and identify any discrepancy.
```

This separates extraction from validation.

8.6 Compare two documents

Attach both documents, then enter:

```
Compare the two attached proposals.

Use the following criteria:
- scope;
- exclusions;
- implementation responsibilities;
- one-time cost;
- recurring cost;
- escalation;
- support;
- data protection;
- termination;
- unresolved assumptions.

For every comparison, cite the relevant file and page.

Do not treat silence as agreement. Write "Not stated" when a proposal does
not address a criterion.
```

8.7 Analyse an image

Claude can examine images supplied in a chat.

Example:

```
Describe this diagram for a beginner.

First state what is visibly present.
Then explain what it may represent.
Keep visible observations separate from interpretation.
Flag any text that is unreadable.
```

This distinction is important:

- **Observation:** “The image contains three boxes connected by arrows.”
- **Interpretation:** “The diagram may represent a three-stage approval process.”

Claude supports image analysis, but conventional photo or illustration generation is a different capability. It can create certain diagrams, charts, HTML, SVG, and interactive visual content through supported features.

[Anthropic: Can Claude produce images?](#)

8.8 Common file-analysis problems

Problem	Likely cause	Correction
Claude misses text	The PDF is a low-quality scan	Upload a searchable PDF or clearer image
Page references are wrong	Page labels and PDF page numbers differ	Ask for both printed page label and PDF page position
A table is corrupted	Complex merged cells or image-based table	Upload the table separately or convert it to spreadsheet format
Claude invents missing terms	Prompt did not restrict the source	State “Use only the attached file” and “Mark missing as not stated”
The file is too large	Context or usage limit	Split it into logical sections
Claude confuses two files	File names are unclear	Rename the files and require file-level citations
A chart is misread	Labels are small or colours are ambiguous	Upload a high-resolution crop and provide relevant context

Knowledge check 8

1. Does a successful upload prove that Claude has read every part correctly?
2. Why should you ask for page or section references?
3. What phrase should Claude use when information is absent?
4. Why should observations be separated from interpretations when analysing an image?
5. What should you do before uploading a confidential work document?

9. Using web search, Research, and extended thinking

Stage purpose

This stage explains how to choose the correct capability when your task depends on current information or deeper analysis.

9.1 When web search is necessary

Use web search for information that may have changed, including:

- Current news.
- Prices.
- Product features.
- Laws and regulations.
- Government guidance.
- Exchange rates.
- Transport schedules.
- Sports results.
- Current office-holders.
- Recent scientific developments.
- Current software documentation.
- Current subscription plans.

Do not rely solely on model memory for these topics.

9.2 Enable web search

The current interface generally uses a tool or sliders control near the prompt box.

1. Open a new chat.
2. Select the sliders or tools control.
3. Locate **Web search**.
4. Turn it on.
5. Enter a search-based prompt.

On Team or Enterprise accounts, an administrator may need to enable the feature. Claude's web-search results include citations or source links that should be reviewed by the user. [Anthropic: Enable and use web search](#)

Visual placement: Insert Visual 5, "Enabling Claude web search," here.

9.3 Use a controlled search prompt

Example:

Use web search.

Explain the current individual Claude plans as of today's date.

Requirements:

- Use Anthropic's official documentation as the primary source.
- State the date checked.

- Distinguish confirmed facts from your interpretation.
- Include plan limits only where the official source states them.
- Cite each material claim close to the relevant sentence.
- Identify anything that may vary by region or account.

Expected output

The response should include:

- Evidence that search was performed.
- Citations.
- A clearly stated date.
- A comparison based on current sources.
- Uncertainty where information is unavailable.

9.4 Inspect the cited sources

Do not stop at seeing a citation symbol.

For an important claim:

1. Open the citation.
2. Confirm that the page is authoritative.
3. Check the publication or update date.
4. Confirm that the source actually supports the claim.
5. Check jurisdiction, product version, population, or scope.
6. Look for conflicting official information.
7. Record the date of access where necessary.

A source can be real but still fail to support Claude's interpretation.

9.5 Web search versus Research versus extended thinking

Capability	Best used for	Example
Ordinary chat	Writing, explanation, supplied-text work	Rewrite an email
Web search	Current, focused factual questions	Find current product documentation
Research	Broader multi-source investigation	Produce a sourced market overview
Extended thinking or higher effort	Difficult reasoning using available evidence	Analyse trade-offs in a complex plan
File analysis	Questions tied to uploaded material	Review a contract draft

Anthropic distinguishes ordinary web search from more extensive Research workflows and from extended thinking. Web search is suited to focused current questions, while Research is intended for broader multi-source investigation. [Anthropic: Choosing Web Search, extended thinking, and Research](#)

9.6 Combine current evidence with reasoning

Example:

```
Use web search to identify the current official requirements.

Then, in a separate section, analyse their implications for a small business.

Keep these categories separate:
1. Confirmed requirements
2. Interpretation
3. Assumptions
4. Recommended actions
5. Items requiring legal or professional confirmation
```

This prevents the analysis from being mistaken for a directly sourced fact.

9.7 Search limitations

Web search can still fail because:

- A source is outdated.
- Search results omit an important document.
- A webpage is inaccessible.
- The relevant information is behind a login.
- Claude misreads a source.
- Two official sources conflict.
- The question uses an incorrect term.
- A cited source supports only part of the answer.
- Information differs by country or user plan.

For consequential matters, compare multiple authoritative sources.

Knowledge check 9

1. Give three examples of questions that require web search.
2. Does the presence of a citation prove that Claude interpreted the source correctly?
3. When is Research more suitable than ordinary web search?
4. What is extended thinking intended to help with?
5. Why should facts and interpretation appear in separate sections?

10. Using Projects, memory, and Artifacts

Stage purpose

This stage introduces features that help organise repeated or multi-step work.

10.1 Create a Project

Projects are available across Claude plans, although limits and advanced knowledge capabilities may differ. Anthropic currently states that free users can create up to five Projects. [Anthropic: What are Projects?](#)

To create one:

1. Open the sidebar.
2. Select **Projects**.
3. Select **Create project** or the equivalent control.
4. Enter a clear project name.
5. Add a short description.
6. Save the Project.

Example name:

| Beginner Cybersecurity Course

| **Visual placement:** Insert Visual 8, "Creating a Claude Project," here.

Expected result: A dedicated Project page should open.

10.2 Add Project instructions

Project instructions tell Claude how to behave across chats inside that Project.

Example:

You are supporting a complete beginner studying cybersecurity.
For all responses:
- use plain English;

- define technical terms when first introduced;
- give an example before advanced theory;
- separate facts from recommendations;
- do not invent sources;
- ask one short knowledge-check question at the end;
- state when current web verification is required.

To add them:

1. Open the Project.
2. Locate the Project instructions area.
3. Paste the instructions.
4. Save them.
5. Start a new chat inside the Project.

Anthropic's Project interface supports adding Project knowledge and instructions. Information from one ordinary chat is not automatically available in another unless it has been placed in the Project's shared knowledge or instructions. [Anthropic: Create and manage Projects](#)

10.3 Add Project knowledge

Project knowledge can include:

- Documents.
- Notes.
- Policies.
- Course material.
- Reference text.
- Code.
- Research material.

To add a document:

1. Open the Project.
2. Locate **Project knowledge**.
3. Select the plus or add-document control.
4. Upload the file or paste text.
5. Give it a clear title.
6. Confirm that it appears in the Project.

Use Project knowledge for relatively stable reference material. Do not treat it as a substitute for current web search when information may have changed.

10.4 Use memory carefully

Where available, memory allows Claude to retain selected preferences or context beyond a single chat.

Potentially useful examples:

- Preferred writing style.
- Standard units or date format.
- Long-term learning objective.
- Repeated formatting preferences.

Do not rely on memory for:

- Critical facts.
- Authoritative records.
- Passwords or secrets.
- Final contractual terms.
- Approved policy versions.
- Exact financial information.

Memory features and their settings may be presented under **Settings -> Memory** or **Settings -> Capabilities**, depending on the current rollout. Users can disable chat search or memory-related capabilities where available. [Anthropic: Chat search and memory](#)

10.5 Create an Artifact

Artifacts can be used to create content that benefits from a separate workspace.

Example prompt:

```
Create an Artifact containing a simple weekly study planner.

Requirements:
- seven rows, one for each day;
- columns for topic, time, activity, and completion;
- no external data;
- accessible headings;
- a simple progress summary;
- no publishing.
```

To enable the required capability where it is not already enabled:

1. Open **Settings**.
2. Open **Capabilities**.
3. Locate **Code execution and file creation**.
4. Enable it after reviewing the permissions.
5. Return to the chat.
6. Request the Artifact.

Artifacts are available on Claude's individual and organisational plans. They may appear in a separate panel and can include documents, diagrams, HTML, SVG, code, or interactive components. [Anthropic: What are Artifacts?](#)

Expected output

A separate panel should open containing the generated planner or interactive output.

10.6 Do not publish sensitive Artifacts

On consumer plans, publishing an Artifact can make it accessible to anyone with the link. Depending on how the Artifact was created, linked conversation material or attachments may create additional exposure risks.

Before publishing:

- Remove personal and confidential information.
- Check attached files.
- Review all generated content.
- Test the Artifact.
- Confirm that no hidden credentials or internal URLs are present.
- Share only with intended recipients.

Anthropic warns that consumer Artifact publication is public-link sharing and that users should be careful when publishing content created from chats, Projects, or attached files. [Anthropic: Publish and share Artifacts](#)

Knowledge check 10

1. What is the main advantage of a Project?
2. Does information from every past chat automatically become Project knowledge?
3. What should Project instructions contain?
4. Why should memory not be treated as an authoritative database?
5. What privacy check is necessary before publishing an Artifact?

11. Practical project: Turn meeting notes into a verified action plan

Project objective

You will use Claude to convert unstructured fictional meeting notes into:

1. A concise meeting summary.
2. A list of confirmed decisions.
3. An action tracker.
4. A budget calculation.
5. A risk and open-question register.
6. A professional follow-up email.
7. A validation report showing whether Claude invented anything.

No programming is required.

Project inputs

Use the following fictional meeting notes:

```
Meeting: Digital Safety Workshop Planning
Meeting date: 2026-08-20
Participants: Lina, Aisha, Omar, and Samir

The team plans to hold a free digital safety workshop for 60 adults.

The proposed event date is 2026-09-18 from 18:00 to 20:00.

The proposed venue is the Community Hub, but the room has not yet been confirmed.

Aisha will confirm the room by 2026-08-25.

Omar will prepare the first draft of the workshop agenda by 2026-08-27.

Marketing must not start until the venue is confirmed. No owner was assigned
for the marketing activity.

The total budget must not exceed AED 4,000.

Current estimates:
- Catering: AED 1,200
- Projector rental: AED 650
- Printing: approximately AED 300

A photo-consent form requires legal review. No owner or deadline was agreed
for this review.

The next check-in will be online on 2026-08-28 at 10:00.
```

Stage 1: Create the workspace

Option A: Regular chat

1. Start a new chat.
2. Name it clearly if chat naming is available.
3. Paste the meeting notes only when requested below.

Option B: Project

1. Create a Project named:

```
Digital Safety Workshop
```

2. Add these Project instructions:

```
Use supplied project material as the source of truth.
```

```
Do not invent owners, dates, decisions, costs, or approvals.  
Distinguish confirmed facts, calculations, inferences, and recommendations.  
Mark absent information as TBD.  
Use ISO date format: YYYY-MM-DD.  
Use AED for monetary values.  
Show formulas for calculations.
```

The Project option is better for repeated work, but the regular-chat option is sufficient.

Stage 2: Perform a basic extraction

Paste the meeting notes and enter:

```
Extract the following from the meeting notes:  
  
1. Meeting purpose  
2. Confirmed decisions  
3. Actions, owners, and deadlines  
4. Costs  
5. Risks  
6. Open questions  
  
Use only the notes.  
  
Do not invent missing information. Mark missing owners or dates as TBD.
```

Expected initial output

Claude should identify, among other things:

- Workshop target: 60 adults.
- Proposed date: 2026-09-18.
- Venue: proposed but not confirmed.
- Aisha's room-confirmation action.
- Omar's agenda action.
- Budget cap: AED 4,000.
- Estimated costs.
- Legal-review owner and deadline: TBD.
- Marketing owner: TBD.
- Next check-in: 2026-08-28 at 10:00 online.

At this point, do not accept the result as final.

Stage 3: Audit the extraction

Enter:

```
Audit your previous answer against the original notes.  
  
Create a table with these columns:  
  
1. Output statement  
2. Exact supporting text from the notes  
3. Confirmed fact, calculation, inference, or recommendation  
4. Supported, partly supported, or unsupported  
5. Correction required  
  
Check especially:  
- owners;  
- deadlines;  
- venue status;  
- marketing status;  
- legal review;  
- costs.  
  
Do not revise the final output yet.
```

Expected output

Claude should identify any unsupported additions.

For example, the following would be errors:

- Assigning Lina as marketing owner.
- Saying that the venue is confirmed.
- Describing the proposed event date as final.
- Giving the legal review a deadline.
- Treating the AED 300 printing estimate as an exact final cost.

Stage 4: Request the budget calculation

Enter:

```
Calculate the total of the stated cost estimates.

Show:
- each amount;
- the formula;
- the estimated total;
- the budget cap;
- the estimated remaining unallocated amount.

Because printing is approximate, label the total and remaining amount as approximate.

Do not add any unstated costs.
```

Expected calculation

Estimated cost:

```
AED 1,200 + AED 650 + AED 300 = approximately AED 2,150
```

Estimated unallocated amount:

```
AED 4,000 - AED 2,150 = approximately AED 1,850
```

This does not mean that the project is definitely under budget. Other costs may not yet have been identified.

Stage 5: Use the complete final prompt configuration

The following is the complete reusable configuration for this project.

```
You are an accuracy-focused meeting assistant.

GOAL
Convert the supplied meeting notes into a verified action package.

SOURCE OF TRUTH
Use only the text inside <meeting_notes>.
Do not use outside knowledge to fill gaps.

RULES
- Do not invent owners, deadlines, approvals, costs, decisions, or event status.
- Mark absent information as TBD.
- Distinguish:
  1. confirmed facts;
  2. calculations;
  3. inferences;
  4. recommendations.
- Use dates in YYYY-MM-DD format.
- Use AED for monetary values.
- Preserve conditional language such as proposed, estimated, and not confirmed.
- Show the formula for every calculation.
- Do not describe the event date or venue as final unless the notes confirm it.
- Keep the follow-up email under 250 words.

OUTPUT FORMAT
A. Executive summary
Maximum 120 words.
```

B. Confirmed decisions

Use a table with:

- decision;
- status;
- supporting evidence.

C. Action tracker

Use a table with:

- action;
- owner;
- deadline;
- dependency;
- status;
- evidence.

D. Financial check

Show:

- item;
- amount;
- formula;
- estimated total;
- budget cap;
- approximate remaining amount;
- limitation.

E. Risks and open questions

Use a table with:

- issue;
- type: risk or open question;
- potential effect;
- owner;
- target date;
- recommended next step.

Do not insert an owner or date when none was agreed.

F. Follow-up email

Include:

- subject line;
- summary of decisions;
- confirmed actions;
- unresolved assignments;
- next check-in.

G. Quality-control report

State:

- whether every action is supported by the notes;
- whether any missing owner or date was marked TBD;
- whether conditional information was preserved;
- whether calculations were independently checked;
- any remaining uncertainty.

<meeting_notes>

Meeting: Digital Safety Workshop Planning

Meeting date: 2026-08-20

Participants: Lina, Aisha, Omar, and Samir

The team plans to hold a free digital safety workshop for 60 adults.

The proposed event date is 2026-09-18 from 18:00 to 20:00.

The proposed venue is the Community Hub, but the room has not yet been confirmed.

Aisha will confirm the room by 2026-08-25.

Omar will prepare the first draft of the workshop agenda by 2026-08-27.

Marketing must not start until the venue is confirmed. No owner was assigned for the marketing activity.

The total budget must not exceed AED 4,000.

Current estimates:

- Catering: AED 1,200
- Projector rental: AED 650
- Printing: approximately AED 300

A photo-consent form requires legal review. No owner or deadline was agreed for this review.

```
The next check-in will be online on 2026-08-28 at 10:00.
</meeting_notes>
```

What the tags mean

The labels `<meeting_notes>` and `</meeting_notes>` mark the beginning and end of the source text.

They help Claude distinguish:

- Instructions.
- Source material.
- Required output.

They are ordinary text markers. No programming knowledge is required.

Stage 6: Review the expected final action tracker

A correct action tracker should resemble the following:

Action	Owner	Deadline	Dependency	Status
Confirm the Community Hub room	Aisha	2026-08-25	None stated	Open
Prepare first agenda draft	Omar	2026-08-27	None stated	Open
Begin marketing	TBD	TBD	Venue must be confirmed	Blocked pending venue
Arrange legal review of photo-consent form	TBD	TBD	Owner and deadline must be assigned	Open
Attend next online check-in	Meeting participants	2026-08-28 10:00	Online access	Scheduled

Important interpretation

“Meeting participants” for the next check-in is a reasonable inference from the notes, not an explicitly assigned action owner. A stricter output may mark the attendee list as “not explicitly stated.” This distinction should be disclosed rather than hidden.

Stage 7: Review the expected risk register

Issue	Type	Potential effect	Owner	Target date	Recommended next step
Venue is not confirmed	Risk	Marketing cannot begin and event planning may be delayed	Aisha for confirmation action	2026-08-25	Confirm or escalate venue booking
Marketing has no owner	Open question	Marketing may not begin promptly after venue confirmation	TBD	TBD	Assign an owner
Legal review has no owner or deadline	Risk and open question	Consent form may not be approved before the event	TBD	TBD	Assign legal-review responsibility
Other costs may be unidentified	Risk	The approximate AED 1,850 balance may be overstated	TBD	TBD	Complete the cost estimate

The recommendation to complete the cost estimate is not a confirmed meeting decision. It must be labelled as a recommendation.

Stage 8: Review the expected email

```
Subject: Digital Safety Workshop - Actions and Open Items

Hello team,

Thank you for the planning meeting on 2026-08-20.
```

The workshop is being planned for 60 adults, with a proposed date of 2026-09-18 from 18:00 to 20:00. The Community Hub is the proposed venue, but the room is not yet confirmed.

Confirmed actions:

- Aisha will confirm the room by 2026-08-25.
- Omar will prepare the first agenda draft by 2026-08-27.

Marketing must not begin until the venue is confirmed. An owner and deadline for marketing remain to be assigned.

The photo-consent form also requires legal review. The owner and deadline for this review remain TBD.

The budget cap is AED 4,000. Current stated estimates total approximately AED 2,150, leaving approximately AED 1,850 unallocated before any additional costs.

Our next check-in is online on 2026-08-28 at 10:00.

Regards,
[Name]

Stage 9: Validate the result

Use this test prompt:

Perform a final validation.

For every owner, date, amount, event status, and decision in the output:

1. show the supporting sentence from the notes;
2. identify whether it is explicit or inferred;
3. flag anything unsupported.

Then recalculate the budget independently.

Do not defend the previous answer. Correct it where necessary.

Validation checklist

The output passes only when:

- Every confirmed action appears in the notes.
- No owner has been invented.
- No deadline has been invented.
- The venue remains "not confirmed."
- The event date remains "proposed."
- Marketing remains conditional.
- Legal review remains unassigned.
- Printing remains approximate.
- The arithmetic equals approximately AED 2,150.
- The approximate remaining amount equals AED 1,850.
- Recommendations are not presented as meeting decisions.

Stage 10: Optional Artifact extension

Enter:

Create a private Artifact containing the verified action tracker.

Requirements:

- display confirmed actions separately from unassigned actions;
- show blocked items clearly;
- include the budget cap and approximate estimated total;
- include a checkbox for completion;
- do not publish the Artifact;
- include no information beyond the supplied meeting notes.

Test the Artifact before relying on it. Do not publish it because this exercise is intended to remain private.

Project knowledge check

1. Why was “proposed” preserved for the event date?
2. Why is the approximate remaining budget not guaranteed to be available?
3. Which two activities have no assigned owner?
4. What event blocks marketing?
5. Why did the validation prompt tell Claude not to defend its previous answer?

12. Visual and screenshot specifications

The following visuals should accompany the tutorial. Screenshots of the actual Claude interface must be captured from a real current account. They must not be generated or reconstructed as fictional software screens.

Visual 1: Illustration - How Claude produces a response

Placement: After Section 1.4, “What a large language model is.”

Purpose: Help beginners understand that Claude generates an answer from a prompt, context, learned patterns, and available tools.

What it should show:

- A user entering a prompt.
- A context area containing earlier messages and optional files.
- A central Claude model represented as an abstract processing system.
- Optional tool inputs such as web search.
- A generated answer.
- A human-review step.

Image-generation instruction:

Create a clean educational vector illustration on a white background. Show a beginner user entering a text prompt on the left. In the centre, show an abstract AI language model made of connected language symbols, without representing a human brain. Above it, show optional context inputs: previous chat, document, and web search. On the right, show a generated response followed by a human review checklist. Label the image clearly as “Conceptual illustration, not the internal technical architecture of Claude.” Use accessible contrast and large readable labels.

Alternative text:

Conceptual flow showing a user prompt and context entering Claude, producing a response that is then reviewed by a person.

Visual 2: Screenshot - Actual Claude sign-in page

Placement: Section 3.1, after opening Claude.

Purpose: Help learners recognise the legitimate starting screen.

What it should show:

- The actual Claude sign-in or registration page.
- Email entry field.
- Available official sign-in options.
- Browser address bar showing the legitimate domain.

Screenshot instruction:

1. Open the official Claude website in a private browser window.
2. Capture the real sign-in page.
3. Include enough of the browser address bar to show the official domain.
4. Do not enter or display a real email address.

5. Crop unrelated browser tabs.
6. Add numbered callouts outside the original interface:
 - 1: Official domain.
 - 2: Email field.
 - 3: Alternative sign-in option.
7. Do not alter the original button labels.
8. Record the screenshot date in the figure caption.

Alternative text:

Actual Claude sign-in page showing the official domain, email field, and available sign-in options.

Visual 3: Screenshot - Main Claude chat interface

Placement: At the beginning of Section 4.

Purpose: Identify the controls used in the tutorial.

What it should show:

- Sidebar.
- New chat control.
- Main conversation area.
- Prompt box.
- Attachment control.
- Tools or web-search control.
- Model selector.
- Send button.

Screenshot instruction:

1. Sign in to a real Claude account.
2. Start a blank new chat.
3. Remove or hide previous sensitive chat titles.
4. Capture the full application window.
5. Redact the user's name, email, and workspace identifiers.
6. Add external numbered callouts without covering interface text.
7. Use the current labels visible in the application.
8. Include a caption stating that interface positions may change.

Alternative text:

Claude chat interface with callouts for the sidebar, new chat, prompt box, attachment control, tools, model selector, and send button.

Visual 4: Illustration - Five-part prompt anatomy

Placement: Section 6.2.

Purpose: Reinforce the G-C-C-F-C prompting framework.

What it should show:

Five connected blocks:

1. Goal.
2. Context.
3. Constraints.
4. Format.
5. Check.

Each block should contain one short question.

Image-generation instruction:

Create a horizontal educational process illustration with five equally sized connected blocks. Use the labels Goal, Context, Constraints, Format, and Check. Under each label include: "What must be done?", "What background matters?", "What rules apply?", "How should it look?", and "How will it be validated?" Clearly label the framework as "Tutorial prompting framework." Use large type, accessible contrast, and no decorative AI robot.

Alternative text:

Five connected prompt components: goal, context, constraints, format, and quality check.

Visual 5: Screenshot - Enabling web search

Placement: Section 9.2.

Purpose: Show how to turn on web search in the current interface.

What it should show:

- A blank chat.
- The tools or sliders control.
- The Web search option.
- Its enabled state.

Screenshot instruction:

1. Open a real blank Claude chat.
2. Select the tools or sliders control.
3. Display the Web search option.
4. Capture the menu with Web search visibly enabled.
5. Do not include private chat history.
6. Add one arrow pointing to the relevant toggle.
7. Do not fabricate menu items or rearrange the interface.
8. State the account plan and screenshot date in the caption because availability can vary.

Alternative text:

Actual Claude tools menu with the Web search option enabled.

Visual 6: Screenshot - Uploading a document

Placement: Section 8.2.

Purpose: Show the attachment process without exposing sensitive data.

What it should show:

- Prompt box.
- Attachment control.
- A harmless sample document shown as attached.
- The file name.

Screenshot instruction:

1. Create a harmless sample file called Workshop_Notes_Sample.txt.
2. Open a blank Claude chat.
3. Select the attachment control.
4. Attach the sample file.
5. Capture the prompt area after the file appears.
6. Ensure no other private files are visible.
7. Add callouts for the attachment control and attached-file indicator.

Alternative text:

Claude prompt box showing the attachment control and an attached sample meeting-notes file.

Visual 7: Screenshot - Privacy and model-improvement settings

Placement: Section 3.3.

Purpose: Help the learner locate consumer privacy controls.

What it should show:

- Settings.
- Privacy section.
- Model-improvement control.
- Explanatory text, where visible.

Screenshot instruction:

1. Open a real Claude account.
2. Navigate to Settings and then Privacy.
3. Capture the section containing the model-improvement control.
4. Redact the account name, email, workspace, and any identifiers.
5. Do not reveal the user's selected preference if privacy is a concern.
6. Add a neutral callout stating, "Review this setting and choose according to your requirements."
7. Include the screenshot date.

Alternative text:

Claude privacy settings showing the control for whether consumer conversations may help improve AI models.

Visual 8: Screenshot - Creating a Project

Placement: Section 10.1.

Purpose: Show the difference between a general chat and a dedicated Project workspace.

What it should show:

- Projects area.
- Create Project control.
- Project name field or newly created Project.
- Project instructions area.
- Project knowledge area.

Screenshot instruction:

1. Open the real Projects interface.
2. Create a harmless sample Project named Beginner Cybersecurity Course.
3. Capture the Project page.
4. Do not upload private documents.
5. Add callouts for Project instructions, Project knowledge, and new chat.
6. Redact account and workspace details.
7. Preserve all genuine interface labels.

Alternative text:

Claude Project workspace showing the Project name, instructions, knowledge area, and new-chat control.

Visual 9: Illustration - Verify before use

Placement: At the beginning of Section 17.

Purpose: Reinforce that consequential AI outputs require review.

What it should show:

A decision path:

- Is the task low consequence?

- Is current or confidential information involved?
- Are authoritative sources available?
- Does an expert need to validate it?
- Use, revise, or stop.

Image-generation instruction:

Create a clear educational decision-tree illustration titled "Verify Before You Act." Begin with "Could an error cause harm, loss, or a wrong decision?" Use branches for low-consequence and high-consequence tasks. For high-consequence tasks, include steps to remove sensitive data, use authoritative sources, independently check calculations and citations, and obtain qualified human approval. Clearly label the image as a conceptual illustration.

Alternative text:

Decision tree showing stronger verification and human approval requirements when an AI error could cause harm or loss.

13. Diagrams and graph decisions

13.1 Why no quantitative graph is included

No quantitative graph is required for this beginner tutorial.

A graph suggesting that:

- Longer prompts always produce better answers.
- More effort always produces more accurate answers.
- A particular model is a fixed percentage better.
- More uploaded data always improves quality.

would create false precision. These relationships depend on the task, evidence, model, plan, prompt quality, and evaluation method.

Workflows and decision trees are more educationally appropriate.

13.2 Diagram: Prompt-improvement cycle

Question answered

How should a learner improve an initial response?




Figure 3. Prompt-improvement cycle.

Interpretation

Do not repeatedly ask Claude to “make it better.” Diagnose the specific problem and provide a targeted correction.

13.3 Diagram: Selecting a Claude capability

Question answered

Should the learner use ordinary chat, files, web search, Research, or a Project?




Figure 4. Selecting a Claude capability.

Interpretation

The correct feature depends on the evidence and duration of the task, not merely on how difficult the prompt appears.

13.4 Diagram: Risk-based review

Question answered

How much verification does a Claude answer require?




Figure 5. Risk-based review.

Interpretation

The higher the consequence of an error, the stronger the evidence and human review must be.

14. Exercises

Exercise 1: Basic explanation

Ask Claude to explain **cloud computing** to a complete beginner.

Requirements:

- Maximum five bullet points.
- No unexplained abbreviations.
- Include one everyday analogy.
- Finish with one knowledge-check question.

Exercise 2: Improve a vague prompt

Rewrite this prompt using G-C-C-F-C:

| Write an email about the delay.

Assume:

- The email is to a software supplier.
- Delivery is 12 days late.
- The agreed delivery date was 2026-08-09.
- You need a revised committed date.
- The tone should be firm but professional.
- The email must not threaten legal action.
- Maximum 180 words.

Exercise 3: Separate facts from assumptions

Give Claude this text:

The training session is planned for October. The meeting room may hold around 40 people. Sarah contacted the facilities team last week. The booking is probably approved because nobody has objected.

Ask Claude to classify each statement as:

- Confirmed fact.
- Estimate.
- Inference.
- Unsupported assumption.

Then ask it to identify what evidence is missing.

Exercise 4: Document-analysis exercise

Use a harmless public or self-created document of two to five pages.

Ask Claude to:

1. Identify the document's purpose.
2. Extract all dates.
3. Extract all stated responsibilities.
4. Provide page references.
5. Mark ambiguous items.
6. Produce a validation table.

Manually check at least three claims.

Exercise 5: Current-information exercise

Enable web search and ask a question about a current software product.

Use this structure:

Use web search to answer this question as of today's date:

[Insert question]

Use only official product documentation for material feature claims.

For each claim:

- cite the source;
- state the relevant version or date;
- identify any plan or regional limitation;
- state what you could not confirm.

The exercise is evaluated on the research method, not merely on the answer.

Exercise 6: Critical-thinking exercise

Choose a proposed decision, such as:

Our company should replace all instructor-led training with AI-generated training.

Ask Claude to provide:

1. The strongest supporting argument.
2. The strongest counterargument.
3. Hidden assumptions.
4. Affected stakeholders.
5. Risks and controls.
6. Conditions under which a pilot would be justified.
7. Measures that would determine success or failure.

Final challenge: Build a Claude learning companion

Create a Project for a subject you want to learn.

The Project must contain:

- A clear name.
- Project instructions for beginner-level teaching.
- At least one approved reference document.
- A seven-day learning plan.
- A daily knowledge check.
- An exercise that increases in difficulty.
- A final assessment.
- A rule against inventing sources.
- A requirement to distinguish evidence from interpretation.
- A privacy review before any file is uploaded.

Final challenge success criteria

Your Project succeeds when:

- Each lesson has a clear objective.
- Technical terms are defined.
- Exercises align with the lesson.
- Answers are supplied separately.
- Current claims use web search.
- Source-based claims refer to the uploaded material.
- Important outputs are checked manually.
- No sensitive information is included.

15. Solutions and model answers

Answers to knowledge check 1

1. AI is the broader field; generative AI specifically creates new content.
2. No. Claude normally generates responses from learned patterns and current context.
3. Examples include summarising documents, drafting emails, explaining topics, and comparing options.
4. Because Claude can produce incomplete, outdated, or incorrect information.

Answers to knowledge check 2

1. The application is the interface; the model is the AI system that generates responses.
2. Context is the information available to Claude for the current response.
3. A hallucination is plausible but false or unsupported generated information.
4. Use web search when information may have changed or requires current evidence.
5. A Project groups related chats, knowledge, and instructions; a normal chat is usually a single conversation.

Answers to knowledge check 3

1. No. Incognito chats may still be retained for a defined period.
2. No. Employer approval and appropriate account controls are still required.
3. In the account's Privacy settings.
4. Recheck the number, wait, and use the retry option rather than requesting many codes rapidly.

Answers to knowledge check 4

1. In the prompt box.
2. No. The default is adequate for most beginner tasks.
3. It provides navigation to chats, Projects, Artifacts, and settings.
4. No. Human validation remains necessary.

Answers to knowledge check 5

1. Claude uses earlier messages in the current chat as context.
2. State the learner level, require definitions, and request shorter sentences.
3. Specify a maximum word, sentence, or bullet count.
4. Generative models may produce different valid wordings across runs.

Answers to knowledge check 6

1. Goal, Context, Constraints, Format, and Check.
2. It tells Claude which evidence should control the answer.
3. No. A role instruction changes perspective, not professional status.
4. When you need consistent structure, tone, or level of detail.
5. To reduce confirmation bias and expose trade-offs.

Answers to knowledge check 7

1. Generate, audit, correct, and finalise.
2. Specific correction identifies the error, correct value, and required recheck.
3. It means connecting a claim to supporting evidence.
4. When the subject changes, instructions conflict, or the conversation becomes cluttered.
5. No.

Answers to knowledge check 8

1. No.
2. They allow you to compare the output with the source.
3. "Not stated," "TBD," or another explicitly defined missing-information label.
4. Because interpretation can add meaning not directly visible in the image.
5. Confirm permission, classification, account approval, and remove unnecessary sensitive content.

Answers to knowledge check 9

1. Current prices, recent laws, present office-holders, current software features, or recent news.
2. No.
3. For broader multi-source investigation.
4. More demanding reasoning and analysis.
5. To avoid presenting the model's interpretation as a sourced fact.

Answers to knowledge check 10

1. It groups related chats, instructions, and reference material.

2. No.
3. Stable rules concerning purpose, style, evidence, terminology, and output.
4. Memory may be incomplete, changeable, or unavailable and is not an authoritative record.
5. Remove sensitive content and inspect files, generated code, links, and sharing scope.

Practical-project answers

1. "Proposed" must remain because the notes do not confirm the event date as final.
2. Other costs may exist, and printing is approximate.
3. Marketing and legal review.
4. Venue confirmation.
5. To encourage a fresh evidence-based audit rather than rationalisation of the earlier output.

Model solution for Exercise 1

Explain cloud computing to someone with no technical background.

Requirements:

- Use no more than five bullet points.
- Define cloud computing in the first bullet.
- Do not use an abbreviation without explaining it.
- Include an analogy based on renting a service rather than owning equipment.
- Finish with one question that checks understanding.

Model solution for Exercise 2

Goal:

Draft an email requesting action on a delayed software delivery.

Context:

The supplier's delivery is 12 days late. The agreed delivery date was 2026-08-09. We require a revised committed delivery date.

Constraints:

- Firm but professional tone.
- Maximum 180 words.
- Do not threaten legal action.
- Do not invent contractual penalties or previous promises.
- Preserve the date exactly.

Format:

- Subject line
- Email body
- Clear request for a revised committed date

Quality check:

After the draft, identify any factual information that should be confirmed before sending.

Example email

Subject: Revised Delivery Date Required

Hello [Name],

The software delivery scheduled for 2026-08-09 remains outstanding and is now 12 days late.

Please provide a clear update on the cause of the delay, the current delivery status, and a revised committed delivery date. We also need confirmation of any outstanding dependency or action required from our side.

The delay is affecting our implementation schedule, so please send the revised plan and confirmed date as soon as possible.

Regards,
[Name]

Model solution for Exercise 3

Statement	Classification	Reason
The training session is planned for October	Confirmed statement, but incomplete date	The text explicitly states it
The room may hold around 40 people	Estimate	"May" and "around" show uncertainty
Sarah contacted facilities last week	Confirmed statement within the text	Explicitly stated, but no documentary evidence supplied
The booking is probably approved	Inference	Derived from lack of objection
Nobody has objected	Claimed fact in the text	It does not establish approval
The booking is approved	Unsupported assumption	No approval evidence is supplied

Missing evidence includes:

- Written room confirmation.
- Exact capacity.
- Date and time of the session.
- Response from facilities.
- Approval reference or booking number.

Model solution for Exercise 6

A strong answer should not simply accept or reject replacing instructor-led training.

It should identify, for example:

- Potential scale and consistency benefits.
- Loss of human facilitation and contextual judgement.
- Accessibility implications.
- Quality assurance.
- Learner engagement.
- Data privacy.
- Assessment integrity.
- Subject suitability.
- Costs of content governance.
- Need for a controlled pilot rather than immediate full replacement.

A reasonable pilot might compare:

- AI-supported learning.
- Instructor-led learning.
- A blended approach.

Measures might include:

- Knowledge retention.
- Completion.
- Learner satisfaction.
- Error rate.
- Accessibility.
- Cost per successful learner.
- Support requests.
- Safety or compliance incidents.

16. Troubleshooting guide

Symptom	Likely explanation	Action
Account creation is unavailable	Unsupported region, network issue, or temporary service problem	Confirm regional support, disable problematic network filtering, and retry later through the official site
SMS code does not arrive	Incorrect number, unsupported number, delay, or repeated attempts	Recheck country code, wait at least five minutes, and select Try again
Claude gives a very generic answer	Prompt lacks context or a specific outcome	Add audience, purpose, constraints, and required format
Claude ignores part of the prompt	Too many conflicting instructions or a long unstructured prompt	Consolidate the requirements and number the steps
Claude gives a false citation	Hallucination or source mismatch	Open the source, verify the claim, and request correction
The answer contains information absent from the file	Claude filled a gap using general patterns	State that the file is the only source and missing information must be marked
Uploaded document is misread	Scanned pages, poor image quality, or complex layout	Upload a searchable version or split the document
A table total is wrong	Extraction or arithmetic error	Ask for formulas, then calculate independently
Web search option is missing	Feature, plan, region, rollout, or administrator restriction	Check account capabilities and organisation settings
Model option is unavailable	Subscription or administrative restriction	Use the default available model
Usage limit is reached	Long chats, files, high effort, or frequent use consumed the allowance	Start a focused chat, shorten context, reduce file size, or wait for usage renewal
Claude keeps following an obsolete instruction	Earlier conversation context is influencing it	Start a new chat with a reviewed handover summary
Artifact does not appear	Required capability may be disabled	Check Settings -> Capabilities -> Code execution and file creation
Artifact behaves incorrectly	Generated code or logic contains an error	Describe the exact failure, ask for a fix, and retest
A published Artifact exposes too much	Sensitive content or attachments were included	Remove publication, sanitise the source, and create a clean version
Response is too short	The prompt does not state the expected depth	Specify sections, examples, and approximate length
Response is too long	The task has no output limit	Set a maximum word count and priority criteria
Claude claims it completed an external action	The required integration may not be connected	Verify the real system, sent message, file, or confirmation rather than trusting the statement
Current information appears outdated	Model memory was used without search	Enable web search and specify the date and official source requirements

Usage limits vary according to plan, model, context length, effort, and enabled capabilities. Long documents and long conversations can consume substantially more capacity. [Anthropic: Usage and length limits](#)

17. Responsible AI considerations

Visual placement: Insert Visual 9, “Verify before use,” here.

17.1 Accuracy and hallucinations

Claude can produce fluent answers that contain:

- Incorrect facts.
- Fabricated quotations.
- Nonexistent references.
- Calculation errors.
- Misread documents.
- Overstated conclusions.

Controls:

- Request sources.
- Open and inspect citations.
- Check primary documents.
- Recalculate important figures.
- Separate fact from inference.
- Ask what Claude could not confirm.
- Obtain expert review for consequential work.

Do not use Claude as the only source for legal, medical, financial, safety, or regulatory decisions. Anthropic explicitly warns that Claude may produce inaccurate or misleading information and should not be treated as a singular authoritative source. [Anthropic: Incorrect or misleading responses](#)

17.2 Privacy and confidentiality

Potential risks include:

- Submitting personal data.
- Uploading confidential contracts.
- Exposing customer information.
- Entering passwords or credentials.
- Publishing sensitive Artifacts.
- Mixing corporate and personal accounts.
- Retaining information in chat history or Projects.

Controls:

- Minimise the data supplied.
- Replace real names with placeholders.
- Remove account numbers and identifiers.
- Use approved organisational accounts.
- Review data-use settings.
- Use Incognito appropriately.
- Understand retention limitations.
- Delete unnecessary chats.
- Do not publish without reviewing every source and attachment.

17.3 Security

AI systems may process content from:

- Uploaded files.
- Webpages.
- External integrations.
- Shared Artifacts.
- Generated code.

Untrusted material may contain misleading instructions intended to influence the AI. Generated code may contain vulnerabilities or unsafe external dependencies.

Controls:

- Do not provide secrets.
- Treat external content as untrusted.
- Inspect generated commands before running them.
- Use test environments.
- Review code and dependencies.
- Restrict external access.
- Use least privilege.
- Avoid importing shared Artifacts from untrusted sources.

Anthropic advises users to review generated files, code, and network activity and to reuse Artifacts only from trusted sources. [Anthropic: Create and edit files with Claude](#)

17.4 Bias and fairness

AI outputs may reflect:

- Imbalances in training data.
- Cultural assumptions.
- Stereotypes.
- Dominant-language perspectives.
- Historical inequalities.
- Bias introduced by the prompt.

Controls:

```
Identify groups or stakeholders who may be disadvantaged by this proposal.  
Check whether the evaluation criteria favour one group without justification.  
Present alternative interpretations from relevant perspectives.  
Do not infer personal characteristics that are not supplied.
```

Human review should include people who understand the affected context.

17.5 Intellectual property and copyright

Potential issues include:

- Uploading material without permission.
- Reproducing protected content.
- Presenting generated text as entirely original human work.
- Using invented references.
- Reusing confidential source material.
- Failing to disclose AI assistance where required.

Controls:

- Use material you are authorised to process.
- Request summaries rather than long reproductions.
- Check institutional or employer rules.
- Verify citations.
- Keep records of AI use where necessary.
- Attribute sources correctly.
- Review originality and licensing before publication.

17.6 High-stakes decisions

Examples include:

- Medical treatment.
- Legal interpretation.
- Recruitment decisions.
- Credit decisions.
- Safety controls.
- Cybersecurity response.
- Financial investment.
- Regulatory compliance.
- Academic assessment.
- Disciplinary action.

Claude may assist with preparation, structuring, or question generation. Final decisions should remain with authorised and qualified humans.

17.7 Overreliance

Frequent AI use can weaken:

- Independent reasoning.
- Writing ability.
- Source evaluation.
- Memory.
- Subject understanding.
- Professional judgement.

A better learning pattern is:

1. Attempt the task yourself.
2. Ask Claude for critique.
3. Compare the two.
4. Investigate disagreements.
5. Produce your own final answer.
6. Explain why you accepted or rejected Claude's suggestions.

17.8 Transparency

Disclose AI assistance when:

- An institution requires it.
- The output affects another person.
- The audience could reasonably assume the work was entirely human.
- AI materially shaped the analysis.
- A professional, academic, or legal policy requires disclosure.

A simple disclosure might state:

Claude was used to help structure the initial draft. The author reviewed the source material, verified the claims, revised the analysis, and remains responsible for the final content.

17.9 Responsible-use control table

Risk	Weak practice	Better control
Hallucination	Accepting fluent text as fact	Verify against authoritative sources
Privacy	Uploading complete confidential files	Minimise and sanitise data

Risk	Weak practice	Better control
Bias	Asking only for support of one view	Request counterarguments and stakeholder impacts
Security	Running generated code immediately	Review and test in a restricted environment
Copyright	Requesting reproduction of an entire protected work	Request analysis or a limited summary
Outdated information	Relying on model memory	Enable web search and state the date
Wrong calculation	Trusting the displayed total	Require formulas and independently recalculate
Overreliance	Letting Claude make the final decision	Preserve qualified human accountability
False attribution	Using references without opening them	Verify source existence and support
Public exposure	Publishing an Artifact without review	Remove sensitive material and inspect sharing scope

18. Summary checklist

Before entering a prompt

- ☐ I know what outcome I need.
- ☐ I have selected the appropriate chat, Project, or tool.
- ☐ I have removed unnecessary personal or confidential information.
- ☐ I have permission to use the supplied material.
- ☐ I know whether current web information is required.

While writing the prompt

- ☐ I stated the goal.
- ☐ I supplied relevant context.
- ☐ I defined important constraints.
- ☐ I specified the output format.
- ☐ I identified the source of truth.
- ☐ I told Claude how to handle missing information.
- ☐ I requested an appropriate quality check.

While reviewing the response

- ☐ The answer addresses the real question.
- ☐ Facts are separated from assumptions and recommendations.
- ☐ Dates, names, amounts, and units are correct.
- ☐ Calculations have been independently checked.
- ☐ Citations exist and support the claims.
- ☐ Current information has been searched and dated.
- ☐ Missing evidence is disclosed.
- ☐ No confidential data has been exposed.

Before using or sharing the result

- ☐ I have edited the output in my own judgement.
- ☐ A qualified person has reviewed high-stakes content.
- ☐ The output complies with relevant organisational rules.

- ☐ Public links and Artifact permissions have been checked.
- ☐ AI assistance is disclosed where required.
- ☐ A human remains accountable for the final decision.

19. Recommended next topics

After completing this tutorial, study the following in order:

1. **Advanced prompt design**
Learn how examples, structured inputs, evaluation criteria, and reusable templates affect consistency.
2. **Document-analysis methods**
Practise evidence tables, page-level traceability, contradiction detection, and source comparison.
3. **Research with Claude**
Learn source selection, search strategies, evidence quality, synthesis, and citation validation.
4. **Projects and knowledge management**
Build repeatable workspaces with controlled instructions and approved sources.
5. **Artifacts and file creation**
Create simple planners, diagrams, reports, and interactive tools while managing publication risk.
6. **AI evaluation**
Develop test cases, scoring criteria, benchmark prompts, and human-review procedures.
7. **Claude API and Claude Code**
Study these only when you are comfortable with basic AI use and, for the API, basic programming concepts.
8. **AI governance and organisational controls**
Learn data classification, permitted-use policies, human oversight, risk assessment, monitoring, and audit evidence.

20. Glossary

Term	Beginner-friendly definition
AI	Computer systems designed to perform tasks associated with intelligence
Anthropic	The company that develops Claude
API	A technical interface that lets software communicate with Claude programmatically
Artifact	A document, application, diagram, or other output created in a separate Claude workspace
Bias	A systematic tendency that may unfairly favour or disadvantage certain views or groups
Chat	A sequence of messages between the user and Claude
Citation	A reference connecting a claim to a source
Claude	Anthropic's family of AI models and assistant applications
Constraint	A rule or limit placed on an output
Context	Information available to Claude while generating a response

Term	Beginner-friendly definition
Context window	The amount of information a model can consider during a request
Consumer account	An individual account intended for personal use rather than centrally managed enterprise use
Data minimisation	Providing only the information necessary for a task
Extended thinking	A setting or capability intended to support more demanding reasoning
Generative AI	AI that creates new content
Hallucination	False or unsupported information generated as though it were true
Human oversight	Review and control performed by a responsible person
Incognito chat	A temporary chat not shown in normal history or used for model training, subject to retention conditions
Inference	A conclusion derived from evidence rather than directly stated by it
Knowledge cutoff	The approximate latest date represented in a model's training knowledge
Large language model	An AI model trained to process and generate language from learned statistical patterns
Memory	A feature that may retain selected context or preferences across conversations
Model	The underlying AI system that processes the prompt and generates a response
Model improvement	Use of eligible data to develop or improve future AI models
Output	The response generated by Claude
Project	A workspace containing related chats, instructions, and reference material
Project knowledge	Documents or text shared within a Claude Project
Prompt	The instruction or information supplied to Claude
Prompt injection	Content designed to manipulate an AI into following unintended instructions
Research	A capability for broader, multi-source investigation
Source of truth	The material that should control the answer
Token	A small unit of text processed by a language model
Traceability	The ability to connect an output statement to supporting evidence
Web search	A tool that lets Claude obtain current information from online sources

21. References

All product-specific details below were checked against Anthropic's official documentation available on 21 August 2026.

-
1. **Get started with Claude:** definition, access methods, supported access requirements, and minimum age. [Anthropic Support](#)
 2. **What can I use Claude for?:** writing, learning, summarisation, translation, coding, brainstorming, and image analysis. [Anthropic Support](#)
 3. **Claude Sonnet 5 and Claude Opus 5 announcements:** current model availability and positioning. [Anthropic Support](#)
 4. **Claude model documentation:** model-specific knowledge-cutoff information. [Anthropic Support](#)
 5. **Account creation and phone verification guidance.** [Anthropic Support](#)
 6. **Prompt engineering guidance:** clear instructions, context, constraints, steps, and examples. [Anthropic Documentation](#)
 7. **Model selection, effort, and extended thinking.** [Anthropic Support](#)
 8. **File-upload guidance and limits.** [Anthropic Support](#)
 9. **Web search and Research guidance.** [Anthropic Support](#)
 10. **Projects and Project knowledge.** [Anthropic Support](#)
 11. **Artifacts, file creation, and publication safety.** [Anthropic Support](#)
 12. **Consumer privacy, model-improvement controls, deletion, and Incognito retention.** [Anthropic Support](#)
 13. **Hallucination and accuracy guidance.** [Anthropic Support](#)
 14. **Usage and context limits.** [Anthropic Support](#)