



A PRACTICAL INTRODUCTION

Everyday AI and Claude for beginners

Understanding large language models — and putting them to work across finance, documents, and communication.

AGENDA

What we'll cover

- 01 What is a large language model
- 02 How a model is trained
- 03 How a response is generated
- 04 Where Claude helps across the business
- 05 Using AI responsibly, and getting started

YOUR PRESENTERS



Dr. Xenia Gleissner
Co-Founder

15 years in cultural heritage. Ph.D. Middle Eastern Studies, Exeter; senior curator across MENA and the UK.



Dev Lakhani
Co-Founder

Two decades in software, cloud and technology, across banking, retail and travel.

What is a large language model


A large language model — an **LLM** — is software that has read an enormous amount of text and learned the patterns within it. It does not store facts in a database or reason like a person. It **predicts likely language**, one **token** at a time, to produce a fluent, relevant response.

TOKEN

The unit a model reads and writes — roughly a word or part of one. "Reconciliation" is a few tokens; "cat" is one. Usage is measured in tokens.

IT PREDICTS THE NEXT WORD

"Please find the invoice"



Learns patterns

Absorbs how words, ideas and structure fit together across billions of examples.



Predicts language

Given your input, it estimates the most fitting next words for the situation.



Responds in plain language

The result reads like natural writing — an answer, a draft, a summary.

How a model is trained

1

Data

A very large collection of text — books, articles and public writing — is gathered as raw material for learning.

2

Training

The model repeatedly tries to predict hidden words, checks itself against the real text, and adjusts — billions of times — until predictions become reliable.

3

Alignment

Human feedback and a published set of principles shape the model to be helpful, honest and safe in how it responds.

HOW IT WORKS

How a response is generated

Your prompt

The question or instruction you type, plus the recent conversation.



Context window

A working memory that holds what's relevant — with a fixed limit.

"I have an interview tomorrow at 3pm." ... the next day ... "How was it?" — Claude knows what "it" means.



Prediction

The model builds the reply one piece at a time, each informed by the last.



Response

A fluent answer, draft or summary returned to you.

INSIDE THE CONTEXT WINDOW

Instructions

Earlier messages

invoice.pdf

Room left

Limit reached → oldest drops off



Clear prompts, better answers

Specific context and instructions steer the prediction toward what you need.



Always verify what matters

It predicts language, so it can be confidently wrong. Check facts and figures.

THE ASSISTANT

Meet Claude

Claude is an AI assistant built by **Anthropic**, designed for safety and judgement and able to “think” through harder problems before answering. You'll meet it across a family of products.

ANTHROPIC'S PRODUCTS



Claude apps

Chat on web, desktop and mobile — the everyday assistant for writing and analysis.



Claude Code

An agent that writes, edits and runs code in your own project — for developers.



Design & artifacts

Build slides, documents and working prototypes you can view and refine in place.



Developer platform

The API and connectors that bring Claude into your own tools and workflows.

Give Claude a solid methodology

Good output comes from a clear method, not a lucky prompt. Agree the ground rules with Claude — then confirm before it acts and review after.



Set the principles

State the standards Claude should apply for this task, and what a good result looks like.



Define what's true

Agree your truth-finding criteria up front, and ask Claude to state the criteria it is using.



Ask it to self-analyse

Have Claude check its own answer against those criteria and flag where it is unsure.



Confirm before, review after

Have Claude restate the task and plan and confirm with you before executing — then review the result against your criteria afterwards.



Use the thinking add-on

For nuanced tasks, turn on extended thinking so Claude reasons step by step before it answers.

Six places Claude earns its keep



Accounting & bookkeeping

Read invoices, categorise expenses, explain the figures in plain English.



Legal terms review

Summarise long contracts and flag the clauses that need a closer human read.



Presentations & reports

Turn rough notes into structured slides and written reports.



Code development

Write, explain and fix code — automation without a developer.



App building

Assemble simple tools and prototypes from a plain description.



Email & communication

Draft and refine messages, adjusting tone and length to fit.

Demo: reconciling the bank feed

The request

March-bank-feed.csv

INV-0421.pdf

INV-0422.pdf

INV-0423.pdf

Match these invoices to the March bank feed by date and amount, allow for invoices raised in a different currency, and flag anything still unpaid.

THINKING · MATCHING ON DATE AND AMOUNT, CONVERTING CURRENCIES AT THE DAY RATE

Matched 2 of 3 invoices — including INV-0422, raised in euros. INV-0423 (£2,100.00) has no matching credit — likely unpaid.

The result

INV-0421

14 Mar · bank credit

£1,240.00

✓ Matched

INV-0422

22 Mar · €1,000.00 → £860.00 at day rate

£860.00

✓ Matched

INV-0423

No matching deposit

£2,100.00

⌚ Unpaid

The bookkeeper reviews and signs off before posting.

Demo: reviewing supplier terms

The request

 Supplier-agreement.pdf

 24 pages

Review these terms and flag anything that affects data ownership or our intellectual property.

THINKING · READING ALL CLAUSES, CHECKING OWNERSHIP AND IP LANGUAGE

Three clauses need attention. Two assign rights away from you — data and IP — and one confidentiality carve-out is unusually broad.

Flagged clauses

Clause 8.2 · Data ownership

Vendor retains rights to all uploaded data

 **Flag**

Clause 11 · Intellectual property

IP in deliverables assigned to supplier, not client

 **Flag**

Clause 14 · Confidentiality

Carve-out broader than standard — worth a look

 **Review**

 A qualified adviser makes the final legal decision.

Going further: skills, memory and agents

Three capabilities move Claude from answering single questions to handling whole tasks — with you reviewing the result.



Skills

Set up your house style, templates and rules once — Claude applies them every time, without being reminded.



Memory

Claude retains useful context across a project or conversation, so you don't repeat yourself each time.



Agents

For multi-step jobs, Claude plans and carries out a sequence of actions — checking in rather than being told every step.

Avoiding “AI slop”

AI slop is the low-value, generic content that appears when people publish AI output without judgement. Easy to spot, corrosive to trust — and entirely avoidable.

What it looks like

- ✗ Bland, padded, obviously templated writing
- ✗ Confident claims that turn out to be wrong
- ✗ The same stock phrases and structure everywhere

How to prevent it



Give real context

Specifics and examples, not a vague prompt.



Always edit

Cut padding; add your voice and facts.



Verify claims

Check every figure, name and fact.



Never publish unread

A first draft is a start, not the end.



Build a skill

Encode your standards — tone, format, sources — so every draft starts in your house style.

Using AI responsibly: privacy, security, carbon



Privacy

Be deliberate about what you enter. Avoid personal or confidential data you're not permitted to share, and know the provider's data terms.



Security

Treat outputs as drafts, keep a person in the loop for consequential decisions, and control who has access.



Carbon

Running large models consumes energy. Use them purposefully, and prefer the right-sized tool for the task.

ANOTHER ROUTE

Open source and self-hostable models

Some models are **open-weight** — you can download and run them on your own hardware, fully offline. Less capable than the leading hosted models, but nothing leaves your building — useful for the most confidential work.



ollama / ollama

Run open models locally with a single command, on Mac, Windows or Linux.

RUN OFFLINE



meta-llama / llama-models

Meta's Llama family — widely-used open-weight models with a permissive licence.

OPEN WEIGHTS



ggml-org / llama.cpp

Runs models efficiently on ordinary laptops and servers — no specialist hardware.

LIGHTWEIGHT



mistralai / mistral-inference

Compact, capable European open-weight models suited to on-premises use.

OPEN WEIGHTS

Sensible pattern: hosted Claude for everyday work, a local model reserved for the most sensitive tasks.

Getting started

01 Start small

Pick one repetitive task — a weekly summary, a recurring email — and try it there first.

02 Be specific

Give context, an example, and the format you want. Clear input produces better output.

03 Keep a human in the loop

Review every result before it's sent, published or filed. You stay accountable.

04 Protect sensitive data

Decide in advance what should never be entered, and set a simple policy for the team.



THANK YOU

Questions and contact

Happy to take questions now. For help applying any of this to your own projects, get in touch.

✉ support@cloudandculture.co