

Lesson 10 — Glossary, cheat sheet & next steps

You made it. Here's the map for going further.

[AI for Dummies](#)[Final Lesson](#)[Reference](#)[Wrap-up](#)

Ten lessons ago you didn't know what an LLM was. Today you can install a local model, call an API, and build a working bot. This slide deck is your permanent reference — vocabulary locked in, cheat sheet packed, road ahead mapped.

Questions you can now answer

Before this course, these were mysteries. Now they're old news.

✓ What's the difference between AI, ML, deep learning, and an LLM?

✓ Why does AI make things up — and what's a hallucination?

✓ What is a token, and why does it matter for cost?

✓ What is a context window and what happens when you exceed it?

✓ How do I run a model on my own computer with no data leaving?

✓ What is an API key and how do I keep it safe?


✓ How do I write a Python script that calls an AI?

✓ What is RAG, and when do I need it?

✓ What is an agent — and what makes it different from a chat?

✓ What is AGI — and does it exist yet?

Glossary spotlight — core concepts



CORE

AI

Software doing tasks that normally need human intelligence — language, images, planning, prediction.

CORE

Machine Learning

AI that learns patterns from data instead of hand-written rules.

CORE

Deep Learning

ML using multi-layer neural networks. Responsible for most AI progress since 2012.

CORE

LLM

Large Language Model: a text-trained model like Claude, ChatGPT, or Gemini. Predicts the next token.

Glossary spotlight — working with AI

The four terms that show up in every single AI conversation.

WORK

Prompt

What you type to the AI.
The whole craft of getting good results is learning to write better prompts.
Recipe: Role + Context + Task + Format.

WORK

Token

A chunk of text — roughly $\frac{3}{4}$ of a word. AI reads and charges by the token, not words or characters. "Hello world" $\approx$ 3 tokens.

WORK

Context window

How much text the model can consider at once. Past the limit, the model "forgets" the beginning of your conversation.

RISK

Hallucination

Confident, false AI output stated as fact, with the same fluent tone as correct answers. You cannot tell from tone alone. **Always verify what matters.**

Remember: confident tone $\neq$ accurate. The model sounds equally sure whether it is right or completely wrong. Verify anything consequential.

Glossary spotlight — doorways & tools

TOOL

API

Application Programming Interface — the way your programs send prompts to and get answers from an AI service over the internet.

TOOL

CLI / Terminal

The text window where you type commands. Used to run Ollama, Python scripts, and Claude Code. Looks intimidating; it's just typing.

TOOL

Local model

A model running on your own computer. No data leaves your machine. Free, private, offline-capable.

WORK

Agent

AI that takes multi-step actions toward a goal — can call tools, browse the web, or run code. Claude Code is an agent.

WORK

RAG

Retrieval-Augmented Generation — feeding the AI snippets of your own documents so it answers from them instead of guessing.

CORE

Multimodal

A model that handles more than text — also images, audio, video, or documents. Most modern frontier models are multimodal.

Glossary spotlight — the bigger picture

BIG PICTURE

AGI

Artificial General Intelligence — hypothetical human-level general AI that can do anything a human can. **Does not exist yet.** Much debated. Don't confuse it with today's impressive-but-narrow systems.

BIG PICTURE

Parameters

A model's internal knobs — “7B” means 7 billion parameters. More parameters generally means smarter but heavier and slower.

BIG PICTURE

Generative AI

AI that creates new content: text, images, audio, video, or code. LLMs are one type; image generators (Stable Diffusion, DALL·E) are another.

Narrow AI is all that exists today — good at specific tasks. AGI (human-level general capability) is a future hypothesis, not a current product.

Model sizes: 3B = phone-friendly • 7B = laptop • 70B = needs a powerful GPU • >100B = datacenter. Start small, scale up as needed.

One-page cheat sheet — part 1

To just use AI

- Go to **claude.ai** and type. That's it.
- Free tier is real — you can do useful work without spending anything.
- Also: ChatGPT (chatgpt.com) • Gemini (gemini.google.com)

To get good answers

- **Recipe:** Role + Context + Task + Format
- Show an example output if you have one
- If the first answer is wrong, say what was wrong and ask again
- Be specific — vague question = vague answer

Example prompt using the recipe: "You are a friendly lawyer [role]. I have a dispute with my landlord about a broken boiler [context]. Explain what rights I have under UK law [task]. Use bullet points, plain English, no jargon [format]."

One-page cheat sheet — part 2

To run AI privately on your computer

- **Easy route (GUI):** Download LM Studio, search for a model, download, run
- **Terminal route:** Install Ollama, then:
`ollama pull llama3.2:3b`
`ollama run llama3.2:3b`
- Nothing leaves your computer. No API key. No costs.

To let AI write & edit code

- Install Claude Code (requires Node.js 18+):
`curl -fsSL https://claude.ai/install.sh |`
`bash`
- Navigate to your project folder and run `claude`
- **Alternative:** Cursor (cursor.sh) — VS Code fork with AI built in

Hardware check: 8 GB RAM runs small models; 16 GB is comfortable; 32 GB+ for the big ones. Apple Silicon Macs are excellent for local AI.

One-page cheat sheet — part 3

To automate with code


- Get an API key at **console.anthropic.com**
- Store it in a `.env` file: `ANTHROPIC_API_KEY=sk-ant-...`
- Install the SDK: `pip install anthropic python-dotenv`
- Start with a cheap model (`claude-haiku-4-5`) & set `max_tokens`
- Only upgrade to a larger model once logic is confirmed working

Always (the three rules)

- **Verify what matters.** AI is a confident guesser, not a fact database.
- **Protect private data.** Don't paste passwords or confidential records into a cloud chat.
- **Keep a human in the loop** for anything medical, legal, financial, or safety-critical.

The full interactive cheat sheet (expandable • printable • 6 sections) lives on the Theory page. Hit Expand all then ⌘P to print a one-pager.

Your 4-week path to confident AI use



This is a practice plan, not a course. Each week adds one capability. The Theory page has a full checkbox tracker — state saved in your browser.

Week 1 spotlight — Chat only

Week
1

Goal: build the habit of actually using AI every day

- Use Claude.ai for at least one real task every day (writing, research, summarising)
- Write a prompt using the full recipe: Role + Context + Task + Format
- Take a mediocre answer and iterate on it until it's good
- Catch the AI making something up — verify with a real source and note the discrepancy
- Explain what an LLM is to someone who doesn't know — without jargon

Don't move on until: you feel comfortable iterating on an answer with follow-up messages, and you've caught at least one hallucination.

Tip: Use it for something you actually care about — a real email, a real question, a real problem at work. Fake exercises don't build habits.

Week 2 spotlight — Go local

Week
2

Goal: run AI on your own machine with no internet

- Install LM Studio (lmstudio.ai) or Ollama (ollama.com)
- Download a small local model (3B–7B parameters)
- Ask the same question to both your local model and Claude.ai — compare quality
- Turn off WiFi and confirm your local model still works

Quick-start commands (Ollama):

```
ollama pull llama3.2:3b  
ollama run llama3.2:3b
```

First run downloads the model (~2 GB). After that it loads in seconds.

Week 3 spotlight — Touch the API

Week
3

Goal: make your first automated AI call from Python

- Get an Anthropic API key at console.anthropic.com
- Store the key in a `.env` file — never in your code
- Run the Hello Claude Python example from Lesson 6
- Process a small batch with a loop (e.g., summarise 5 items)
- Check your usage console — understand what you spent

```
# The whole thing in 8 lines
import anthropic
client = anthropic.Anthropic()           # reads ANTHROPIC_API_KEY from env
msg = client.messages.create(
    model="claude-haiku-4-5",
    max_tokens=256,
    messages=[{"role": "user", "content": "Say hello in one sentence."}],
```

```
)  
print(msg.content[0].text)
```

Week 4 spotlight — Build one thing

Week
4

Goal: ship something real, however small

- Pick one starter project from Lesson 8 (or write your own idea)
- Get a "hello world" working — even if it's ugly
- Make it do the one thing you actually wanted it to do
- Show it to one other person (even if they aren't technical)

Starter project ideas

- Summarise all your email newsletters in one daily digest
- AI that answers questions about your own notes / documents
- Script that rewrites your job descriptions in plain language

The mindset shift: You're not trying to build the perfect thing. You're trying to go from zero-to-shipped once — because the second project is always easier than the first.

- CLI tool that translates files in bulk
-

Where to go next — curated links

[Anthropic Prompting Guide](#) ↗

Official techniques, examples, and patterns for prompting Claude.

[Anthropic API Docs](#) ↗

Complete reference for the Anthropic API — endpoints, SDKs, examples.

[Ollama](#) ↗

Download & run local models via terminal. One command per model.

[LM Studio](#) ↗

Desktop app for local AI models — no terminal, friendly GUI.

[Hugging Face](#) ↗

The GitHub of AI models. Browse thousands, run demos in-browser.

[Andrej Karpathy \(YouTube\)](#) ↗

Deep but accessible explanations of how neural networks & LLMs work.

[Cursor](#) ↗

VS Code fork with AI built in — great if you prefer an editor over a terminal tool.

[fast.ai](#) ↗

Free practical deep learning course — best path from beginner to practitioner.

The one rule — reprised

//

AI is confidently wrong sometimes.

Verify what matters.

What “verify” looks like: if the AI tells you a legal deadline, a drug interaction, a financial figure, or any claim you’re going to act on — check it against a primary source. The AI is not that source.

What it does NOT mean: you don’t need to fact-check every sentence. Use judgment. Verify the things that matter; let the draft writing slide.

This habit — more than any technical skill — is what separates a confident, responsible AI user from someone who gets burned. Build it in Week 1 and never lose it.

You finished AI for Dummies

//

You don't need to understand how an engine works to drive a car. The same is true here. Start with Chat, build the verify-everything habit, and let curiosity pull you toward the next thing. That's it.

Lesson 10 complete

Course complete

You're dangerous now

Go build something

Next step: open the [Theory page](#), bookmark it, and start the 4-week tracker.

Then: open [Homework](#), pick your final project, and write your AI elevator pitch.

