



# How to Use This AI Course

Welcome | The 16 folders and what each one owns, four reading paths (practitioner, ML engineer, researcher, leader), how every sheet is built, a realistic pace and a spaced-review schedule

[#StartHere](#) [#AI](#) [#StudyPlan](#) [#ReadingPaths](#)

## What this course is

### This AI course



A stack of short, dense sheets that take you from school maths to the AI systems of 2026, every term explained where it is used.



A reference you can open on any sheet: formulas with their variables defined, numeric worked examples, tables of the tools and their trade-offs.

- each sheet stands alone
- concepts live in one folder only



A map of the field with the vocabulary, the costs and the risks you need to talk to teams and vendors, plus the governance rules that apply.

#### PICTURE IT

Like a set of laminated cards from an experienced colleague: one card per topic, nothing on it you would not use.

The same course described to three readers. Pick the sentence that sounds like you and follow the matching reading path below

16

folders

plus a generated glossary

2-5

pages per sheet

longer topics are split in parts

≥ 8

flashcards per sheet

answers printed

2026

state of the art

reviewed date on every sheet

### Every answer is on the page

No exercise in this course leaves you a blank to fill. Quizzes, flashcards and self-assessments print the answer and the reason next to the question: cover the answer with a card, answer aloud, then check. Reading the answer without trying first is the slow way.

## The 16 folders and what each owns

Each notion is taught in **one** folder only. Other sheets recall it in one sentence and send you to its home. The table gives the home of each field, in reading order.

| Folder      | What you learn there                                                                                                | Read after |
|-------------|---------------------------------------------------------------------------------------------------------------------|------------|
| 00-welcome  | how to study, big picture, roadmap, vocabulary, your workbench                                                      | nothing    |
| maths       | notation, linear algebra, calculus, optimisation maths, probability, statistics, information theory, floating point | welcome    |
| programming | Python, NumPy, pandas, PyTorch, JAX, ML software engineering, CUDA basics, SQL                                      | welcome    |

|                           |                                                                                                         |                      |
|---------------------------|---------------------------------------------------------------------------------------------------------|----------------------|
| algorithms                | complexity, data structures, graphs, dynamic programming, search, nearest-neighbour algorithms          | programming          |
| data                      | data lifecycle, quality, features, labelling, web corpora, synthetic data, formats                      | programming          |
| machine-learning          | framing, classic models, metrics, validation, clustering, time series, interpretability                 | maths, data          |
| deep-learning             | neural networks, backprop, optimisers, CNNs, attention, transformers, diffusion, training at scale      | machine-learning     |
| llm                       | tokens, embeddings, pretraining, scaling laws, fine-tuning, RLHF, RAG, serving, evaluation, models 2026 | deep-learning        |
| agents                    | agent loop, tool use, MCP, memory, planning, multi-agent, coding and computer-use agents                | llm                  |
| reinforcement-learning    | MDPs, Q-learning, policy gradients, PPO, model-based RL, RL for LLMs                                    | maths, deep-learning |
| perception-and-generation | vision, image and video generation, speech, audio, 3D, robotics, classic NLP                            | deep-learning        |
| mlops                     | system design, tracking, pipelines, serving, monitoring, LLMOps, cost and latency                       | machine-learning     |
| compute-infrastructure    | GPUs and accelerators, interconnects, memory and FLOPs, maths, data centres, schedulers                 | deep-learning        |
| cloud                     | AWS, Azure and Google AI services, model APIs, AI certifications                                        | mlops                |
| responsible-ai            | safety, interpretability, fairness, privacy, AI security, regulation, governance                        | llm                  |
| culture                   | history, landmark papers, labs, industry, society, careers, how to stay current                         | any time             |

### Folders have their own start sheet

Every folder except this one opens with 00-start-here (scope, prerequisites, reading order inside the folder) and ends with a self-assessment. In 00-welcome, this sheet plays the start-here role and there is no self-assessment.

© Four reading paths

| Tool                      | Practitioner | ML engineer | Researcher | Leader |
|---------------------------|--------------|-------------|------------|--------|
| 00-welcome                | ✓            | ✓           | ✓          | ✓      |
| maths                     | ⚠            | ✓           | ✓          | ✗      |
| programming               | ✓            | ✓           | ✓          | ✗      |
| algorithms                | ✗            | ⚠           | ⚠          | ✗      |
| data                      | ⚠            | ✓           | ⚠          | ⚠      |
| machine-learning          | ⚠            | ✓           | ✓          | ⚠      |
| deep-learning             | ⚠            | ✓           | ✓          | ✗      |
| llm                       | ✓            | ✓           | ✓          | ⚠      |
| agents                    | ✓            | ✓           | ⚠          | ⚠      |
| reinforcement-learning    | ✗            | ⚠           | ✓          | ✗      |
| perception-and-generation | ⚠            | ⚠           | ✓          | ✗      |
| mlops                     | ✓            | ✓           | ✗          | ⚠      |
| compute-infrastructure    | ⚠            | ✓           | ⚠          | ⚠      |
| cloud                     | ✓            | ✓           | ✗          | ⚠      |
| responsible-ai            | ✓            | ✓           | ✓          | ✓      |
| culture                   | ⚠            | ⚠           | ✓          | ✓      |

✓ yes / full    ⚠ partial    ✗ no / none

Which folders each path reads: full = read every sheet, partial = read the start-here and the sheets named in the table below, no = skip on the first pass


| Path         | Goal                                      | Partial folders: which sheets                                                                                                                                 | First-pass size    |
|--------------|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| Practitioner | build products on models and APIs         | maths: notation, vectors, probability; ML: framing, metrics; DL: attention, transformer; data: lifecycle, synthetic data; compute: memory maths               | approx. 90 sheets  |
| ML engineer  | train, serve and run models in production | algorithms: complexity, hashing, ANN; RL: RL for LLMs; perception: vision tasks, speech; culture: papers, careers                                             | approx. 160 sheets |
| Researcher   | read and write papers, run experiments    | algorithms: DP, search; data: corpora, synthetic; compute: accelerators, memory maths                                                                         | approx. 150 sheets |
| Leader       | decide, budget, govern, hire              | data: lifecycle; ML: framing, causal inference; llm: landscape, evaluation; agents: what an agent is; mlops: cost; compute: data centres; cloud: services map | approx. 50 sheets  |

- 1 Read 00-welcome first sheets 02 and 04 give the words every later sheet uses
- 2 Open a folder at 00-start-here check its prerequisites before going on
- 3 Read the sheets in the .order sequence easy and useful first, specialised last
- 4 Close the folder with its self-assessment below 80 % correct: revisit the sheets it points to

5 Come back for depth on the second pass the partial folders become full when your work needs them

How to walk any path

How every sheet is built



The anatomy of a sheet. The same blocks appear in the same roles on every sheet, so you always know where to look

**Formula block** — Every formula lists each symbol with its meaning and unit, then a worked substitution; computed visuals (memory budgets, scaling laws, sampling) calculate their numbers from the inputs shown.

**Callouts** — Tip = a practical shortcut; warning = a common mistake with consequences; real world = a dated case; definition = a term to learn by heart.

**Reviewed date** — The day the sheet was last checked against the state of the art; AI moves fast, so trust mechanisms over model names and numbers older than a year.

**As of 2026 / approx.** — Marks a figure that changes quickly or is an order of magnitude; never a precise price or benchmark score.

Before Start the AI course, can you explain...

|                                                                                                                                                                                                | Not yet                  | Partly                   | Yes                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|--------------------------|--------------------------|
| <div>MUST KNOW</div> <div>Read technical English</div> <div>the course is in English; French official terms are glossed; Learn: any tech blog, with a dictionary</div>                         | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| <div>MUST KNOW</div> <div>School arithmetic: fractions, powers, percentages</div> <div><math>10^9</math>, <math>2^{10} \approx 1,000</math>, 25 % of 80; Learn: maths: algebra refresher</div> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| <div>MUST KNOW</div> <div>Use a computer and a terminal</div> <div>open files, install an app, type a command; Learn: sheet 05 of this folder</div>                                            | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| <div>HELPFUL</div> <div>Basic Python</div> <div>variables, loops, functions; Learn: programming: Python for AI</div>                                                                           | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| <div>HELPFUL</div> <div>Secondary-school algebra and graphs</div> <div>functions, slopes, logs; Learn: maths folder</div>                                                                      | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| <div>ADVANCED</div> <div>Calculus and linear algebra</div> <div>only for ML engineer and researcher paths; Learn: maths folder</div>                                                           | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |


What the course assumes before sheet 02. Anything marked must: an evening of revision is enough

🕒 Pace and spaced review

|                                     | Mon    | Tue    | Wed    | Thu    | Fri    | Sat | Sun    |       |
|-------------------------------------|--------|--------|--------|--------|--------|-----|--------|-------|
| Read one new sheet (first pass)     | 45 min | •      | 45 min | •      | •      | •   | •      | 2×/wk |
| Redo the worked example by hand     | •      | 30 min | •      | 30 min | •      | •   | •      | 2×/wk |
| Flashcards of due sheets            | 10 min | 10 min | 10 min | 10 min | 10 min | •   | •      | 5×/wk |
| Hands-on: notebook or small project | •      | •      | •      | •      | •      | 2 h | •      | 1×/wk |
| Self-assessment or mixed review     | •      | •      | •      | •      | •      | •   | 40 min | 1×/wk |

A sustainable week for a working adult: about 6 hours, of which one third is recall and practice, not reading

Sheet: AI big picture (02)



| # | Review      | Day | Date       | Done | Score / notes |
|---|-------------|-----|------------|------|---------------|
| 0 | First study | 0   | Mon 5 Oct  | ✓    |               |
| 1 | Review 1    | 1   | Tue 6 Oct  | ✓    | 5 / 9         |
| 2 | Review 2    | 3   | Thu 8 Oct  | ✓    | 7 / 9         |
| 3 | Review 3    | 7   | Mon 12 Oct | ✓    | 8 / 9         |
| 4 | Review 4    | 14  | Mon 19 Oct | ✓    | 9 / 9         |
| 5 | Review 5    | 30  | Wed 4 Nov  | ✓    | 9 / 9         |

Example schedule for one sheet first read on 5 October 2026, intervals 1-3-7-14-30 days. The notes show a typical learner's flashcard scores rising across reviews

Y

How many weeks does the practitioner path take at 6 hours per week?  
Given: Practitioner first pass: approx. 90 sheets, 3.5 pages each on average; First reading: approx. 12 minutes per page; recall and practice add 50 % on top; Study time available: 6 hours per week

Y

Step 1: How many pages is that?

90 sheets × 3.5 pages = 315 pages  
the page is the unit you can time

T

Y

Step 2: How long is the first reading?

315 pages × 12 min = 3,780 min = 63 h  
dense sheets read slower than prose

T

Y

Step 3: Add recall and practice?

63 h × 1.5 = 94.5 h  
flashcards, worked examples redone, small projects

T

Y

Step 4: Convert to weeks?

94.5 h ÷ 6 h/week ≈ 16 weeks  
divide by the weekly budget

T

About 16 weeks, roughly four months, for a practical command of the practitioner path  
Check: At 10 hours per week:  $94.5 \div 10 \approx 9.5$  weeks; doubling study time halves the calendar, not the effort

T

Turning a reading path into a calendar (author's estimates for a first pass)

Study rules that the research on learning supports

✓ baseline / mandatory    ☐ recommended – verify in your context

- ✓ Test yourself before re-reading: cover the answer, answer aloud, then check retrieval practice
- ✓ Review on a schedule (1, 3, 7, 14, 30 days), not when you feel like it spaced review
- ✓ Redo every worked example on paper without looking
- ☐ SHOULD Mix two or three folders in a week once you are past the welcome folder interleaving
- ☐ SHOULD Build something small each month: a notebook, a script, a tiny app
- ☐ COULD Explain one sheet to someone else in five minutes gaps show up immediately

MYTH OR FACT?  
“Highlighting and re-reading make it stick”

MYTH: Answering questions does  
Re-reading builds familiarity, which feels like knowing.  
Retrieval practice (testing) produces more durable memory in  
study after study.

MYTH OR FACT?  
“AI moves too fast to learn from sheets”

MYTH: Mechanisms last; product names do not  
Gradient descent (1847), backpropagation (1986) and attention  
(2014-2017) are still the core. Learn mechanisms here; follow  
news for names.

MYTH OR FACT?  
“Read every folder in order before building anything”

MYTH: Alternate reading and building from the first  
month  
Sheet 05 sets up your workbench so that each folder can be  
tried in a notebook the same week.

Three study habits that feel productive and are not

Remember

TOPIC  
How to use this course

The five lines to keep from this sheet.

HEADLINE 1  
One notion, one home folder

Detail: other sheets recall it in a line

HEADLINE 2  
Pick a path, then go folder by folder

Detail: start-here first, self-assessment last

HEADLINE 3  
Every sheet has the same anatomy

Detail: entry point, prerequisites, core, worked example,  
myths, takeaways, flashcards

HEADLINE 4  
Six hours a week, one third recall

Detail: approx. four months for the practitioner path

HEADLINE 5  
Review at 1, 3, 7, 14, 30 days

Detail: test before you re-read

Self-test

- Q1

Where is the attention mechanism taught, and where are the LLM block tweaks (RoPE, GQA)?

A

Attention in deep-learning; modern LLM block tweaks in llm. Each notion has one home folder.
- Q2

What opens and what closes every folder except 00-welcome?

A

00-start-here opens it (scope, prerequisites, order); a self-assessment with answers closes it.
- Q3

Which two folders does every reading path read in full?

A

00-welcome and responsible-ai.
- Q4

Which path skips maths on the first pass, and why can it?

A

The leader path: it needs vocabulary, costs and risks, not derivations.
- Q5

Give the five spaced-review intervals used in this course.

A

1, 3, 7, 14 and 30 days after the first reading.
- Q6

What does 'approx.' or 'as of 2026' mean on a sheet?

A

An order of magnitude or a fast-changing fact; never a precise price or benchmark.
- Q7

What is retrieval practice and why prefer it to re-reading?

A

Answering from memory before checking; it builds durable memory, re-reading only familiarity.
- Q8

Roughly how long is the practitioner first pass at 6 h/week?

A

About 16 weeks: approx.  $315 \text{ pages} \times 12 \text{ min} \times 1.5 \approx 95 \text{ h}$ .
- Q9

Below what self-assessment score should you revisit a folder?

A

Below 80 % correct: reread the sheets the missed questions point to.