

AI for Everyone

No maths, no code, no jargon. What AI actually is, what it can and cannot do, and how to use it well. Written for students, teachers, and anyone who keeps hearing about it.



Utsav Poudel
Founder, Fulcrum

Who made this, and why it is free

Fulcrum is a volunteer-run non-profit. We help people from under-resourced regions learn artificial intelligence through one-on-one mentorship, research and publication supervision, project support, workshops, and a guest speaker series.

The premise is simple. Ability is spread evenly across the world; opportunity is not. What separates a capable student in Pokhara or rural Bihar from one at a well-funded lab is rarely intelligence. It is access to supervision, to reviewers, to equipment, and to the unwritten rules nobody writes down.

This deck is part of how we try to close that. It is free, it will stay free, and you may reuse and adapt it for your own teaching under CC BY 4.0. If it is useful to you, come and find us.

Find us

Website

fulcrumgo.github.io

More material

fulcrumgo.github.io/resources

Discord

discord.gg/gbQCGkupdJ

LinkedIn

linkedin.com/company/gofulcrum

Every deck we publish is free to download at
fulcrumgo.github.io/resources

PART ONE

It is a pattern machine, not a mind

Almost every confusion about AI comes from getting this one thing backwards.

What AI actually is

Strip away the language and it is simpler than it sounds.

- ▲ A program that finds **patterns in examples**, then applies them to something new.
- ▲ Show it thousands of photos labelled 'cat' and it learns what tends to be true of cat photos. Nobody writes the rules down.
- ▲ It does not understand cats. It has a very good statistical sense of what cat-shaped data looks like.
- ▲ That is the whole trick, repeated at enormous scale.

This is why AI is brilliant at things with lots of examples, and unreliable at things it has rarely seen.

A useful test

Ask: could this be learned from examples? If yes, AI probably does it well. If it needs judgement about something genuinely new, be careful.

Why it arrived so suddenly

- ▲ The ideas are decades old. What changed recently is **data** and **computing power**.
- ▲ The internet produced enormous quantities of text and images to learn from.
- ▲ Graphics chips, built for games, turned out to be ideal for training these models.
- ▲ Once both existed, methods that had always been theoretically sound started working in practice.

So it feels like a sudden breakthrough. It is closer to a long slow build that finally crossed a threshold.

What a chatbot is really doing

- ▲ It predicts **the next word**, over and over, based on everything before it.
- ▲ It was trained on a vast amount of text, so its predictions are usually sensible and fluent.
- ▲ It is optimised for **plausibility**, not truth. Those two usually coincide, and sometimes do not.
- ▲ When it invents a fact, a citation, or a quote, that is the same machinery working normally, not a glitch.

It is a superb writer with an unreliable memory. Treat it that way and you will rarely be caught out.

What it is genuinely good at

- ▲ **Explaining** something you find difficult, at whatever level you ask for.
- ▲ **Rewriting and summarising** text you already have.
- ▲ **Translating** between languages, and between technical and plain wording.
- ▲ **Getting you started**, a first draft, an outline, an idea to react to.
- ▲ **Answering embarrassing questions** patiently, at two in the morning, without judgement.

The pattern

It excels when you supply the material and it transforms it. It is weakest when you ask it to supply facts from memory.

Where it goes wrong

- ▲ **Confident invention.** It will produce a fake reference in the same tone as a real one.
- ▲ **Out-of-date knowledge.** It knows what it was trained on, not what happened last week.
- ▲ **Inherited bias.** Trained on human text, it reproduces human prejudice.
- ▲ **No sense of its own limits.** It cannot reliably tell you when it does not know.

Rule: if it matters, verify it somewhere else. Especially numbers, citations, medical and legal claims.

PART TWO

Using it to study, without it doing the studying

The question students ask most: doesn't this make us worse at thinking?

“Doesn’t using AI make us dumb?”

It can. It depends entirely on which of two things you do with it.

- ▲ **Learning with it**, asking it to explain, to quiz you, to check your reasoning. This makes you better.
- ▲ **Outsourcing to it**, having it produce the answer you then hand in. This makes you worse, quietly.
- ▲ The danger is that outsourcing feels productive. Work gets finished, and nothing is learned.
- ▲ The honest test: **could you now do it without the AI?** If not, you did not learn it.

Learn with AI. Do not let AI learn for you. That single sentence is most of the guidance you need.

Try this instead

Attempt the problem first. Then ask AI to critique your attempt. You keep the thinking and still get the help.

Asking well

- ▲ **Give context.** 'Explain recursion to a 15-year-old who knows loops' beats 'explain recursion'.
- ▲ **Say what you want back**, a list, a paragraph, a worked example, a counter-argument.
- ▲ **Ask it to show its reasoning**, then check the reasoning rather than trusting the answer.
- ▲ **Push back.** 'What is wrong with that?' and 'What would a critic say?' produce far better results than accepting the first reply.

Tools worth knowing

- ▲ **Chat assistants**, explanation, drafting, and rewriting.
- ▲ **Wolfram Alpha**, maths and science, computed rather than predicted, so the answers are actually reliable.
- ▲ **Translation tools**, good enough for meaning, still imperfect for nuance.
- ▲ **Search with AI summaries**, quick orientation, but follow the links before you cite anything.

For anything numerical, prefer a tool that calculates over one that predicts. This distinction matters more than any brand.

Jobs, honestly

- ▲ Tasks change faster than whole jobs disappear. Most roles absorb the tools and move on.
- ▲ The people affected first are those doing routine work that is easy to describe in text.
- ▲ The people who gain most are those who understand their own field **and** know how to use the tools.
- ▲ You do not need to become an engineer. You need to stop being intimidated by it.

Where to go next

- ▲ Use one assistant properly for a month on real work. Practical familiarity beats theory here.
- ▲ If you enjoy it, learn a little **Python**. It is the language nearly all of this is built in.
- ▲ Our **Machine Learning: Foundations** deck is the natural next step, and assumes no prior knowledge.
- ▲ Bring a question to our Discord. Genuinely, no question is too basic there.

ABOUT FULCRUM

This was written for you to keep

Share it, teach from it, print it for a classroom. If you want to go further, we mentor people through exactly that, free.

Fulcrum is a volunteer-run non-profit helping people from under-resourced regions learn AI through free mentorship, research supervision, and open material like this deck. Founded 2025 in Kathmandu, Nepal. Nothing we do costs anything.

Website

fulcrumgo.github.io

Every deck, free

fulcrumgo.github.io/resources

Apply for mentorship

discord.gg/gbQCGkupdJ

Partnerships

linkedin.com/company/gofulcrum