

Activity Guide

AI Foundations.

Learn by doing: ask an AI three questions, judge three answers, and catch a made-up fact

TIME

About 25 minutes

LEVEL

Beginner, no tech background

WORKS ON

Phone, tablet, or computer

GOES WITH

Course 1 video, Lessons 1–6

About 25 minutes. Free ChatGPT or Claude works.

Learner record

Name _____

Date _____ Quiz score _____ / 10

Cohort or self-paced _____



BEFORE YOU BEGIN

About this activity

Watching tells you how AI works. Trying it shows you. In this guide you put the course into practice: you **explain the six big ideas in your own words**, then **test a real AI tool** with three kinds of questions and judge each answer, and finally **go hunting for a hallucination** and check it with the three quick checks from Lesson 5.

BY THE END YOU WILL HAVE

- Explained AI, training, tokens, and hallucinations in your own words
- Rated three real AI answers and written how you checked each one
- Caught (or tried to catch) a made-up fact, and proven it
- Ticked off everything the quiz will ask about

WHAT YOU NEED

- A phone, tablet, or computer with internet
- A free ChatGPT or Claude account (the free plans are enough; having both is a bonus)
- A pen, or a copy of this PDF you can type into
- About 25 minutes

STAY SAFE WHILE YOU PRACTICE

Don't type personal details into an AI chat: no full name with address, ID numbers, passwords, bank details, or private information about other people. Everything in this activity works without them.

How this guide works

Part	What you do	Time
A · Remember	Six short questions, one per lesson, in your own words	5 min
B · Try	Ask an AI tool three questions, rate the answers, record how you checked	10 min
C · Hunt	Ask for something AI often makes up, then check it	5 min
D · Ready	A checklist before you take the quiz	2 min
Answer guide	What good answers look like (look after you try)	3 min



PART A · REMEMBER · ABOUT 5 MINUTES

The six big ideas, in your words

Answer in one or two sentences, without looking back at the slides. If you get stuck, rewatch the lesson named in brackets, then try again. Writing it yourself is what makes it stick.

1. How is AI different from regular software? Use the recipe idea if it helps. (Lesson 1)

2. A spam filter was never given a rule for a brand-new scam, yet it caught it. How? (Lesson 2)

3. Why might an AI trained mostly on photos of male doctors make unfair guesses? (Lesson 2)

4. In one sentence, how does a chatbot write its answer? Use the words *next word* or *token*. (Lesson 3)

5. Name two things AI is good at and two things it struggles with. (Lesson 4)

6. What is a hallucination, and what are the three quick checks? (Lessons 5 and 6)

Three questions, three answers

Open ChatGPT or Claude. Ask one question of each type below, either the example or one of your own. Read each answer carefully, give it a rating, and write down how you checked it.

1

QUESTION 1 · A FACT

Ask something with one right answer

Facts are easy to check, so they show you how accurate the tool is.

EXAMPLE TO TYPE

How tall is the Eiffel Tower, in meters?

Your own idea: a capital city, the year something was invented, the boiling point of water.

My question	Its answer (short)	Rating	How I checked
		☐ Correct ☐ Partly ☐ Wrong	☐ Source ☐ Second opinion ☐ Common sense

Tip: the tower's height changes slightly over the years as antennas are added, so a trusted source may give a newer number. That's a real-world lesson in checking dates, too.

2

QUESTION 2 · AN OPINION

Ask something with many good answers

There's no single right answer here, so you judge whether it's sensible and useful.

EXAMPLE TO TYPE

What's the best way to study for an exam?

Your own idea: how to start a hobby, how to plan a small party, how to save time in the morning.

My question	Its answer (short)	Rating	How I checked
		☐ Helpful ☐ So-so ☐ Not helpful	☐ Source ☐ Second opinion ☐ Common sense

- ☐ **a** Now ask the same question again, in a **new chat**. Is the second answer worded differently? (Lesson 3: answers vary on purpose.)

3

QUESTION 3 · SOMETHING RECENT

Ask about this week

This tests the tool's training cut-off, and whether it admits what it doesn't know.

EXAMPLE TO TYPE

What happened in the news this week?

My question	Its answer (short)	Rating	How I checked
		☐ Correct ☐ Partly ☐ Wrong	☐ Source ☐ Second opinion ☐ Common sense
What did it do?			Tick one
Said it can't know recent events, or that its knowledge stops at a date			☐
Searched the web and showed sources you could open			☐
Answered confidently with no sources			☐ (check carefully!)

☐ **Checkpoint.** All three rows are filled in, and you asked Question 2 twice.

C

PART C · HUNT · ABOUT 5 MINUTES

Catch a hallucination

Hallucinations hide in quotes, citations, statistics, and very specific details (Lesson 5). Ask for one, then try to prove it's real.

TYPE THIS

Give me a quote about curiosity from a famous scientist, with the exact book title and page number.

- ☐ **a** Copy the quote, the person, and the book it gave you into the table.
- ☐ **b** **Source check:** search the web for the exact quote in quotation marks, and for the book title. Can you find the book? Does the quote appear in it, or on a trusted quotes site that names where it's from?
- ☐ **c** **Second opinion:** ask the other tool, or the same one in a new chat, "Is this a real quote? Where is it from?"
- ☐ **d** **Common sense:** does the book exist, and was it written when that person was alive?

What it said	What I found
Quote	
Person	
Book and page	

My verdict: ☐ Real and checked ☐ Partly real (right person, wrong book or page) ☐ Made up ☐ Couldn't tell

DIDN'T FIND A HALLUCINATION?

Good tools are getting better at this, and many now search before answering. That's still a pass: you checked instead of trusting. Try once more with something obscure, such as a quote from a less famous person, or the exact population of a small town.

What will you do differently the next time an AI gives you a fact you plan to use?

D

PART D · READY · ABOUT 2 MINUTES

Before you take the quiz

I can...

☐

Explain how AI differs from regular software (rules written by people vs patterns learned from examples)

☐

Put AI, machine learning, and generative AI in order, from biggest to smallest

☐

Describe training vs using, and why the data decides what an AI learns

☐

Explain next-word prediction, tokens, and why answers vary

☐

List what AI is strong at, and six things it struggles with

☐

Spot a hallucination, and use source, second opinion, and common sense

☐

SUBMIT IT, THEN TAKE THE QUIZ

Go to aicafe.dev/courses/ai-foundations and log in. Under **Activity**, write a short summary of your three questions, ratings, and checks. Then take the quiz: ten questions, eight to pass. Missed it? Review, wait ten minutes, and try again. Your certificate of completion is issued as soon as both are done.



LOOK AFTER YOU TRY

Answer guide

Your words will differ. These show what a strong answer includes.

#	A strong answer includes
A1	Regular software follows rules a person wrote (a recipe). AI learns its own rules from many examples (tasting thousands of dishes).
A2	It learned patterns from thousands of labeled emails (words like FREE!!!, strange links, unknown senders) and recognized them in the new email. It gives a likelihood, not a certainty.
A3	It copies patterns in its data. If most examples show men, it may assume doctors are men. Garbage in, garbage out: it doesn't know what's fair.
A4	It predicts the most likely next token (a word or part of a word), adds it, and repeats until the answer is done.
A5	Good at: drafting, summarizing, explaining, brainstorming, translating, coding help. Struggles with: made-up facts, exact math, recent news, knowing your situation, judgment calls, and it sounds sure even when wrong.
A6	A confident but false answer, common in quotes, citations, and statistics. Checks: source, second opinion, common sense.
B1	Any answer you confirmed on a trusted site. If the numbers differ slightly, note which is newer.
B2	A sensible answer you'd actually try. The second answer should be worded differently, with similar ideas.
B3	Best: it says it can't know, or searches and shows sources. Riskiest: a confident answer with no sources.
C	Any verdict you backed up with what you found. "Couldn't tell" is honest; "Real" with no source check is not a pass.