



Keep Your Brain

Initial Learner Diagnostic

Purpose: Learn how the learner currently uses AI, how they reason about its outputs, where they already demonstrate AI-literacy habits, and what interests can be used as entry points for future sessions.

Target time: ~20 minutes

This is **not a test of AI vocabulary**. The goal is to observe behavior, reasoning, habits, and current mental models.

Do not teach every concept as it appears. Ask, observe, and follow up.

The learner should leave feeling like:

“They wanted to know what I already know and what I want to do.”

—not:

“I just took an AI exam.”

Tutor Stance

Ask before explaining.

Let the learner disagree with the AI and with you when they can explain their reasoning.

Do not correct every misconception immediately. A misconception discovered during intake may become useful evidence for deciding what to teach later.

Do not require disclosure of private experiences.

When asking about real AI use, the learner may describe the situation without showing private chats, school records, messages, images, account information, or identifying details.

Record **observable behavior**, not personality judgments.

1. Current AI Use

Approx. 4 minutes

Start casually:

Before we actually start doing stuff, I want to figure out how you already use AI. There aren't really right answers here. I'm mostly trying to make sure I don't spend time teaching you things you've already figured out.

Ask

1. What AI stuff have you used before?



Examples only if needed:

- ChatGPT
- Gemini
- Copilot
- Snapchat or social-media AI
- image generators
- AI features inside games or apps
- something else

Notes:

2. What do you usually use AI for?

Notes:

3. Tell me about the last time you used AI for something you actually cared about.

Listen for:

- goal
- why AI was chosen
- how the learner prompted it
- whether the learner revised the request
- what they did with the result

Notes:

4. What's something you've tried to get AI to do where it just completely sucked?

Do not supply explanations first.

Listen for how the learner explains failure:

- "It was wrong."
- "It didn't understand me."
- "I had to explain it differently."
- "It made something up."
- "It ignored part of what I asked."
- "The conversation got confused."
- "I didn't give it enough information."
- another explanation



Follow naturally with:

What did you do when that happened?

Notes:

2. How the Learner Works With AI

Approx. 6 minutes

These questions probe habits rather than terminology.

There is no need for the learner to know words such as *hallucination*, *provenance*, *cognitive ownership*, or *loaded input*.

Trust

When AI gives you an answer, how do you decide whether it's right?

Follow-up if useful:

Does that change depending on what you're asking it?

Listen for whether trust comes from:

- confidence
- detail
- familiarity
- agreement
- citations
- outside evidence
- existing knowledge
- independent verification

Notes:

Recovery

If AI gives you a bad answer, what do you usually do next?

Listen for:

- regenerate
- repeat the prompt
- rewrite the prompt
- add context



- correct the model
- challenge assumptions
- ask for alternatives
- start a fresh chat
- check elsewhere
- abandon AI

Notes:

Error Recognition

Have you ever caught AI making something up or getting something wrong? How did you figure it out?

Listen for the difference between:

“It seemed wrong.”

and

“I checked.”

Notes:

Verification

Say AI gives you three sources for something. Do you trust the answer more because it has sources? What would you actually check?

Listen for:

- whether the source exists
- who produced it
- date
- original/primary source
- whether the learner would open it
- whether the source actually supports the claim
- whether several sources are independent

Notes:



Data Boundaries

Is there anything you wouldn't put into an AI chat? Why?

Possible areas the learner may independently identify:

- passwords
- contact information
- private messages
- school records
- health information
- identifying information
- someone else's information
- photos or likenesses
- confidential material

Do not turn missing examples into a quiz.

Notes:

Cognitive Ownership

If AI could just do your homework for you, when would you want it to give you the answer, and when would you want it to help you figure it out?

Listen for whether the learner distinguishes:

finishing the task

from

building the ability.

Notes:

3. Live AI Challenge

Approx. 7 minutes

The purpose is to observe the learner directing AI rather than asking them to describe how good they are at prompting.

Give the learner this scenario:

Pretend you have a science test Friday and you don't understand photosynthesis very well. You have ChatGPT. Show me what you'd type.



Do not coach the first attempt.

Record the learner's first prompt or describe it:

Observe

Does the learner:

- state a goal?
- provide useful context?
- specify how they want help?
- choose a role for the AI?
- ask for an explanation?
- ask for examples?
- ask for questions or a quiz?
- provide their current understanding?
- ask the AI to expose gaps?
- ask for sources?
- create checkpoints?
- simply ask for the answer?

Do not treat a long prompt as automatically better than a short one.

Prompt Revision

After the first attempt, ask:

Now change your prompt so the AI helps you actually learn photosynthesis instead of just giving you a good explanation.

Let the learner decide what to change.

Record the revised approach:

Follow-up

Ask:

What changed between your first prompt and your second one?

Then:

How would you know whether what it taught you was actually correct?

Optional final probe:



How could you tell whether you actually learned it instead of just getting a good answer?

Listen for:

- explaining from memory
- answering questions
- applying the idea
- checking a textbook or reliable source
- correcting an error
- producing an original example
- doing a similar task without AI

4. Find the Learner's Door

Approx. 3 minutes

End by moving away from assessment and toward possibility.

Ask:

Forget school for a second. If I could show you how to use AI to make, learn, investigate, or build basically anything, what sounds fun?

Let the learner answer freely first.

If they need ideas:

Game? Story or fictional world? Art? Video? Coding? Learning something you're obsessed with? Investigating weird stuff? Practicing conversations? Inventing something? Something completely different?

Do not require a STEM interest.

Learner interests

Things they would like AI to help them do

Possible future project or demonstration



Tutor Observation Record

Complete this **after** the conversation when possible.

Do not make the learner sit through the rubric.

Use the existing Keep Your Brain progression:

Emerging

The learner notices the issue after it is pointed out and needs support choosing the next action.

Practicing

The learner predicts some behavior, recognizes the concept when it appears, and selects an appropriate response with light support.

Independent

The learner recognizes the pattern in a new context and responds without being led through every step.

Transfer-ready

The learner combines several concepts, explains tradeoffs, and can apply or teach the reasoning in another context.

1. AI Use

What does the learner currently use AI for?

This is descriptive rather than scored.

2. Direction

Can the learner deliberately change what the AI does?

Look for:

- useful context
- changing instructions after failure
- assigning a role
- setting constraints
- recognizing when a fresh approach is needed

Current evidence:

Observed level: _____



3. Skepticism

Does the learner distinguish sounding right from being right?

Look for:

- questioning confident outputs
- recognizing unsupported specificity
- noticing assumptions
- willingness to say "I don't know"
- distinguishing agreement from evidence

Current evidence:

Observed level: _____

4. Verification

Does the learner have a method beyond asking the AI again?

Look for:

- opening sources
- checking dates
- identifying source type
- comparing independent evidence
- checking whether evidence actually supports the claim
- appropriately lowering confidence

Current evidence:

Observed level: _____

5. Ownership

Does AI support the learner's thinking or replace it?

Look for:

- making their own decisions
- attempting work before requesting completion
- asking for hints, criticism, questions, or practice
- explaining results without the chat
- distinguishing their contribution from AI's contribution

Current evidence:



Observed level: _____

6. Boundaries

Does the learner recognize when information or a task should not enter an AI system?

Look for:

- personal information
- information belonging to someone else
- consent
- school or environmental rules
- high-stakes situations
- knowing when a person is the appropriate next step

Current evidence:

Observed level: _____

Interest and Motivation

Do **not** treat this as proficiency.

This determines which door into the curriculum is most likely to matter to this learner.

Strongest interests

Things that produced visible curiosity

Things that produced little interest

Learner-generated project ideas

Best current entry point

Planning the Next Session

Based on the diagnostic:



Habits already demonstrated independently

Habits worth strengthening

Largest current gap

Interest to use as the teaching context

Keep Your Brain concepts most relevant next

Concepts that may not require much instructional time

Pilot Evidence

For each learner, preserve the same basic evidence:

Initial use: What are they already doing with AI?

Initial prediction/reasoning: What do they currently expect AI to do?

Observed behavior: What did they actually do during the live challenge?

Verification behavior: How did they propose checking an output?

Human decision: What choice did they retain for themselves?

Interest: What did they want to explore?

Transfer evidence later: What can they eventually do in a new situation without step-by-step prompting?



Facilitator Reminder

The diagnostic exists to answer:

Where is this learner starting?

How do they currently think about AI?

What can they already do without help?

Where does their judgment break down?

What do they actually want to do with these tools?

The curriculum does not need to start at the same place, move at the same speed, or use the same examples for every learner.

Keep the literacy goals stable.

Change the door.

And throughout the program:

Predict → Run → Notice → Name → Decide → Transfer