BrightKit

AI Literacy Assessment

Pre / Post instrument · Version 1.0

A 10-item instrument measuring four constructs of AI literacy in children ages 7–14: AI mental model, hallucination detection, bias awareness, and verification behavior.

Designed to be administered twice: once before the BrightKit AI Safety curriculum (pretest) and once after (posttest).

INSTRUMENT DETAILS

Items:	10 (8 multiple choice, 3 open response)
Constructs:	M1, M2, M3, M4
Age range:	7–14 (grades 3–8)
Administration:	Paper or digital
Version:	1.0 · April 2026
Licence:	Creative Commons BY 4.0

Administered as part of the BrightKit AI Safety curriculum study
brightkit.app/research · arun@voltexam.com

Student Information

Fill out before you begin (your answers are anonymous)

Student ID (assigned by your teacher): _____

Grade level: _____

Age: _____

Today's date: _____

Is this your FIRST time taking this assessment, or your SECOND?

■ First (pretest) ■ Second (posttest)

Before you start:

- There are 10 questions. Circle the letter of the answer you choose, or write in the space provided for open questions.
- Answer honestly. This is not graded. Your teacher only sees totals.
- If you don't know an answer, it's okay to write 'I don't know' or 'I'm not sure.'
- Take your time. There is no time limit.

Assessment — Page 1 of 5

Circle your answer. Take your time.

M1 · MENTAL MODEL

Q1. When an AI chatbot answers a question, which of these best describes what it is doing?

- ○ a) It is thinking the way a person thinks.
- ○ b) It is looking up the answer in a big book.
- ○ c) It is guessing the most likely next word, based on patterns it was trained on.
- ○ d) It is calling a human on the other side of the screen.

M1 · MENTAL MODEL

Q2. In your own words: what do you think AI actually is? (1–2 sentences)

M2 · HALLUCINATION DETECTION

Q3. An AI says: 'George Washington invented the light bulb in 1782.' What is most likely true?

- ○ a) The AI is right — I learned something new.
- ○ b) The AI is wrong, but it sounds sure.
- ○ c) Both Washington and Edison invented the light bulb.
- ○ d) I can't tell, so I'll just believe it.

M2 · HALLUCINATION DETECTION

Q4. Which of these is MOST likely a sign an AI has made something up?

- ○ a) It answers very quickly.
- ○ b) It gives you a confident answer about a book, person, or event — but the details don't check out anywhere else.
- ○ c) It uses long words.
- ○ d) It adds emoji to the answer.

M3 · BIAS AWARENESS

Q5. An AI picture-maker always draws a 'doctor' as an older white man. Why is this most likely happening?

- ○ a) The AI is being mean on purpose.
- ○ b) The AI was trained on pictures where most doctors were shown as older white men.
- ○ c) Most doctors in real life are older white men.
- ○ d) AI can only draw in one style.

M3 · BIAS AWARENESS

Q6. If an AI doesn't work well in your grandmother's language, what is the MOST likely reason?

- ○ a) Her language is too hard.
- ○ b) The AI wasn't trained on enough examples of her language.
- ○ c) AI doesn't like old people.
- ○ d) AI only speaks English on purpose.

M4 · VERIFICATION BEHAVIOR

Q7. You ask an AI a science question for your homework. It gives you a confident answer. What do you do next?

- ○ a) Write it down — the AI said so.
- ○ b) Check the answer in at least one other place (a book, a trusted website, a teacher).
- ○ c) Ask a different AI the same question and use whichever sounds smarter.
- ○ d) Skip the question.

M4 · VERIFICATION BEHAVIOR

Q8. An AI tells you a fact. You ask where it got the information from. The AI's source doesn't exist when you search for it. This is a signal that:

- ○ a) The search engine is broken.
- ○ b) The AI may have made the source up. The original answer might not be reliable.
- ○ c) The source is just very old.
- ○ d) The AI is hiding something from you.

M4 · VERIFICATION BEHAVIOR

Q9. Describe ONE time you would NOT use AI to decide something. Why? (1–2 sentences)

M4 · VERIFICATION BEHAVIOR — SCENARIO

Q10. An AI helps you write most of an essay for class. Your teacher asks if you wrote it yourself. What would you do, and why? (2–3 sentences)

For Teachers — Scoring Guide

Do not distribute to students. Keep separate from completed assessments.

Multiple-choice answer key

Q1 (Mental Model):	c — pattern prediction
Q3 (Hallucination):	b — confidently wrong
Q4 (Hallucination):	b — details don't verify
Q5 (Bias):	b — training data
Q6 (Bias):	b — underrepresented training data
Q7 (Verification):	b — cross-check
Q8 (Verification):	b — fabricated source

Open-response scoring (Q2, Q9, Q10)

Q2 — What is AI?

2 pts: Student describes AI as pattern-based, statistical, or learning from data. Example: 'AI is a computer that learns patterns from lots of examples.'

1 pt: Student identifies AI as a tool or computer program without the pattern/data mechanism. Example: 'AI is a computer that answers questions.'

0 pts: Student treats AI as a thinking being, magic, or a person. Example: 'AI is a smart robot brain.'

Q9 — When would you NOT use AI to decide?

2 pts: Student names a high-stakes or personal-judgment context AND gives a reason connected to AI's limits (mistakes, bias, missing context). Example: 'I wouldn't use AI to pick a friend — only I know what feels right.'

1 pt: Student names a reasonable context OR gives a reason, but not both clearly.

0 pts: Student says 'I always use AI' or can't identify any context.

Q10 — The essay scenario

2 pts: Student identifies honesty as the right action AND articulates a reason rooted in integrity, learning, or fairness to classmates. Example: 'I would tell my teacher AI helped — because writing is how I learn, and it's not fair to say I did it all myself.'

1 pt: Student identifies honesty as correct but gives a surface reason ('because lying is bad').

0 pts: Student says they would lie or wouldn't tell the teacher, with no reasoning about consequence.

For Teachers — Score Interpretation

How to read pre/post results

Scoring totals

Multiple choice (Q1, Q3–Q8): 1 point per correct answer. Maximum: 7 points.

Open response (Q2, Q9, Q10): 0–2 points per question using the rubric above. Maximum: 6 points.

Total possible: 13 points.

Interpreting pre/post change

We recommend comparing each student's posttest score to their own pretest score, not against an external benchmark. Meaningful gain thresholds:

• **+1 to +2 points:** modest shift, likely reflects improved mental model.

• **+3 to +5 points:** substantial gain across two or more constructs. Typical after the 5-lesson curriculum.

• **+6 or more:** strong gain. Worth noting students who reach this level — they often become peer leaders in subsequent classes.

What to share, and with whom

Totals by class, not by student. Share with your principal, parents at a curriculum night, or your district AI advisory (if one exists). Anonymous class averages are suitable for external communication.

Individual scores should stay in the classroom and are most useful for identifying students who would benefit from extension discussions — not for grading.

Reporting back to BrightKit (optional)

If your school is running the pilot as part of the formal study, we ask you to share anonymous class-level totals. No individual student data is ever requested. The data helps refine the curriculum and strengthens the research base. Opt in by emailing arun@voltexam.com.