



TEACHING MAP

AI Literacy Teaching Map



A teachable sequence for learners roughly ages 12–15

KEEP YOUR BRAIN • USE THE MACHINE



Companion resource: AI Literacy Prompt Labs (AI_Literacy_Prompt_Labs.md) contains three copy/paste-ready prompts for every concept in this map.

Purpose

This map turns the 16-part AI literacy framework into a teachable sequence for learners roughly ages 12–15.

The core learning method is **side-by-side comparison**. The learner runs two or three deliberately different prompts, observes what changes, explains why it changed, and decides what to do with the result.

The prompts are not the lesson by themselves. Each comparison has a specific job:

1. Make one hidden AI behavior visible.
2. Give the learner language for what happened.
3. Practice a response that preserves judgment and agency.
4. Check whether the learner can transfer the habit to a new situation.

The sequence moves from noticing influence, to verifying outputs, to managing safety and power.

Learning Loop

Use the same six-part loop for every concept.

1. Predict

Before running anything, ask the learner what they expect each prompt to produce.

This prevents the activity from becoming passive entertainment.

2. Run

Run two or three prompts in fresh chats when possible. Keep the model and settings the same so the changed prompt is the main variable.

3. Notice

Compare specific features rather than asking which answer is “better.” Depending on the lesson, inspect:

- claims
- evidence
- tone
- assumptions
- uncertainty
- sources



- requested actions
- private information
- responsibility

4. Name

Give the observed behavior a durable name: loaded input, false certainty, copied claim, unsafe disclosure, outsourcing, or another concept from the framework.

5. Decide

The learner chooses what to trust, reject, verify, revise, disclose, or escalate.

6. Transfer

Give a new example without announcing the concept. Mastery means the learner recognizes the pattern and responds without being led through every step.



Phase 1: Mind Games and Influence

The first phase teaches the learner to notice how wording, tone, and presentation shape belief before evidence has been checked.

1. AI Is a Yes-Man

LEARNING OUTCOME

The learner can distinguish emotional agreement from useful judgment and can deliberately ask an AI to challenge their point of view.

COMPARISON DESIGN

- Prompt A presents a weak idea with confidence and asks for support.
- Prompt B presents the same idea neutrally and asks for an evaluation.
- Prompt C asks for the strongest disagreement or missing perspective.

WHAT SHOULD BECOME VISIBLE

- The AI may mirror the user's confidence and framing.
- Warmth and agreement can make a weak answer feel trustworthy.
- A different role or instruction can produce a meaningful challenge.

WHAT TO COMPARE

- How quickly each answer agrees
- Whether risks or counterarguments appear
- How the tone affects the learner's reaction
- Whether the underlying evidence changed

MASTERY EVIDENCE

The learner identifies agreeable language without treating it as proof and can write a follow-up that invites genuine challenge.

TRANSFER CHECK

Give the learner a flattering response to an unrelated bad idea. Ask what they would do next.

2. Loaded Inputs

LEARNING OUTCOME

The learner can find an assumption hidden inside a question and rewrite the question so more than one conclusion remains possible.

COMPARISON DESIGN

- Prompt A contains an accusation or conclusion inside the question.



- Prompt B describes the same facts without assigning blame.
- Prompt C asks what information would be needed before reaching a conclusion.

WHAT SHOULD BECOME VISIBLE

- A question can steer an answer before the model begins responding.
- AI often continues a false or incomplete premise.
- Neutral wording changes both the tone and the range of possible answers.

WHAT TO COMPARE

- Assumptions accepted by each answer
- Missing information requested
- Range of explanations considered
- Strength of the conclusion

MASTERY EVIDENCE

The learner underlines the hidden claim and rewrites the prompt in neutral language.

TRANSFER CHECK

Ask the learner to repair a loaded question about a school rule, friendship conflict, or news event.

3. Fluency Is Not Accuracy

LEARNING OUTCOME

The learner treats polished language as presentation rather than proof and marks details that require verification.

COMPARISON DESIGN

- Prompt A asks for a confident, detailed answer without requiring sources.
- Prompt B asks for the same answer with uncertainty clearly marked.
- Prompt C restricts the answer to facts that can be supported by supplied material.

WHAT SHOULD BECOME VISIBLE

- Detail and confidence can exist without factual support.
- An answer may mix accurate information with one invented specific detail.
- Constraints can make uncertainty easier to see.

WHAT TO COMPARE

- Exact names, numbers, quotations, dates, and links
- Claims presented without support
- Statements of uncertainty
- Whether polish changes perceived credibility

MASTERY EVIDENCE

The learner flags checkable claims before repeating them and explains how each could be verified.

**TRANSFER CHECK**

Give the learner a polished paragraph containing one suspiciously specific claim.

4. Confidence Is Not Evidence

LEARNING OUTCOME

The learner separates a claim, the tone used to present it, and the evidence that would support it.

COMPARISON DESIGN

- Prompt A asks for a forceful argument.
- Prompt B asks for evidence supporting and challenging the same claim.
- Prompt C asks the AI to label every statement as claim, evidence, inference, or uncertainty.

WHAT SHOULD BECOME VISIBLE

- A confident paragraph can add no new evidence.
- Repetition and formatting can create a false sense of support.
- Evidence has a traceable relationship to a specific claim.

WHAT TO COMPARE

- Number of claims versus pieces of evidence
- Relevance of each source or example
- Unsupported leaps
- Language that implies certainty

MASTERY EVIDENCE

The learner can point to the evidence for a claim—or state that evidence is missing.

TRANSFER CHECK

Show a confident recommendation and ask what evidence would change the decision.



Phase 2: Verification and Cognitive Ownership

The second phase turns skepticism into a working practice. The learner checks outputs, keeps the final decision, and uses AI to increase ability rather than replace it.

5. Verification Skill

LEARNING OUTCOME

The learner can trace an important claim to an original source and explain exactly what that source does and does not support.

COMPARISON DESIGN

- Prompt A asks for an answer with citations.
- Prompt B asks for primary sources, dates, and relevant excerpts or locations.
- Prompt C asks for a verification plan without making the factual claim itself.

WHAT SHOULD BECOME VISIBLE

- A citation can exist without supporting the sentence attached to it.
- Multiple pages may repeat one original source.
- Verification requires opening and inspecting the source.

WHAT TO COMPARE

- Source type and date
- Whether the source exists
- Whether it supports the whole claim
- Independence of the sources

MASTERY EVIDENCE

The learner produces a short verification note: claim, source, support, limitation, and confidence.

TRANSFER CHECK

Give a new cited claim and ask the learner to verify it without a checklist.

6. Keep the Wheel

LEARNING OUTCOME

The learner treats AI recommendations as inputs and explains the human values, facts, rules, and risks behind a final decision.

COMPARISON DESIGN

- Prompt A asks the AI to make the decision.
- Prompt B asks for several options with tradeoffs.



- Prompt C states the learner's priorities and asks which facts remain unknown.

WHAT SHOULD BECOME VISIBLE

- A recommendation depends on goals and context the AI may not know.
- Different priorities can produce different reasonable choices.
- The final decision remains with the person affected.

WHAT TO COMPARE

- Assumptions about the learner's goals
- Missing local or personal context
- Tradeoffs acknowledged
- Who is presented as responsible

MASTERY EVIDENCE

The learner accepts, revises, or rejects a recommendation and gives a reason independent of "the AI said so."

TRANSFER CHECK

Use a low-stakes choice in a different domain and ask the learner to explain their own decision rule.

7. Spar, Don't Outsource

LEARNING OUTCOME

The learner can redesign an AI request so it builds skill rather than bypassing the work.

COMPARISON DESIGN

- Prompt A asks for the finished answer.
- Prompt B asks for one hint or one question at a time.
- Prompt C provides the learner's attempt and asks for critique without rewriting it.

WHAT SHOULD BECOME VISIBLE

- Finishing a task and learning the skill are different outcomes.
- A tutor-style interaction requires the learner to keep producing work.
- Feedback becomes more useful when there is a real attempt to inspect.

WHAT TO COMPARE

- Who performs the difficult step
- What the learner must think or produce
- Whether understanding can be tested afterward
- How much useful struggle remains

MASTERY EVIDENCE

The learner converts a "do it for me" prompt into a tutor, critic, quiz, or practice prompt.

TRANSFER CHECK

Give a new assignment request and ask the learner to create an AI interaction that increases independent ability.



8. Defend What You Submit

LEARNING OUTCOME

The learner can explain an AI-assisted result, identify the contribution of AI, and reproduce the important reasoning without the chat.

COMPARISON DESIGN

- Prompt A produces a polished result immediately.
- Prompt B requires the AI to ask for the learner's reasoning before helping.
- Prompt C asks the AI to conduct a short oral defense after the result is finished.

WHAT SHOULD BECOME VISIBLE

- A polished answer can exceed the learner's actual understanding.
- Follow-up questions reveal missing ownership quickly.
- Explanation is evidence of understanding.

WHAT TO COMPARE

- Ability to explain vocabulary and choices
- Ability to recreate key steps
- Ability to identify and correct an error
- Clarity about which parts came from AI

MASTERY EVIDENCE

The learner gives a concise defense covering the main idea, process, source, personal choice, and limitation.

TRANSFER CHECK

Remove the chat and ask the learner to teach the result to someone else.



Phase 3: Digital Boundaries, Safety, and Power Use

The third phase applies judgment to privacy, synthetic media, scams, consent, rules, disclosure, and deliberate high-leverage use.

9. Data Boundaries

LEARNING OUTCOME

The learner decides what information can be shared, what must be changed, and when the tool should not be used.

COMPARISON DESIGN

- Prompt A includes unnecessary identifying details.
- Prompt B replaces identifying details with neutral placeholders.
- Prompt C asks for a way to complete the task without uploading the sensitive material.

WHAT SHOULD BECOME VISIBLE

- Useful prompts often need less personal information than expected.
- Removing a name may not remove identity.
- Sometimes the safe choice is not to use AI for the task.

WHAT TO COMPARE

- Information necessary for the task
- Direct and indirect identifiers
- Information belonging to someone else
- Safer alternatives

MASTERY EVIDENCE

The learner can classify details as necessary, unnecessary, sensitive, or prohibited and can create a safer version.

TRANSFER CHECK

Present a new message, image, or school record and ask what could safely enter a chatbot.

10. Synthetic Evidence

LEARNING OUTCOME

The learner treats images, audio, and video as claims requiring provenance and context rather than self-proving evidence.



COMPARISON DESIGN

- Prompt A asks the AI to judge authenticity from appearance alone.
- Prompt B asks for several possible explanations of the media.
- Prompt C asks for a provenance-checking plan using source, date, location, and independent reporting.

WHAT SHOULD BECOME VISIBLE

- Visual inspection cannot reliably prove authenticity.
- Real media can be paired with a false caption.
- Provenance checks are stronger than confidence about appearance.

WHAT TO COMPARE

- Certainty of the conclusion
- Alternative explanations considered
- Evidence requested outside the media itself
- Willingness to wait before sharing

MASTERY EVIDENCE

The learner creates and follows a provenance checklist before reacting or sharing.

TRANSFER CHECK

Give a plausible image with an unsupported caption and ask what is known, unknown, and checkable.

11. Scam Mechanics

LEARNING OUTCOME

The learner recognizes the pressure pattern—urgency, high emotion, fake authority, and immediate action—and slows down.

COMPARISON DESIGN

- Prompt A asks the AI to respond immediately to a high-pressure message.
- Prompt B asks the AI to identify pressure tactics before considering the request.
- Prompt C asks for an independent verification route that does not use the message's links or contact details.

WHAT SHOULD BECOME VISIBLE

- Pressure tries to reduce the time available for judgment.
- The requested action matters more than polished wording.
- Independent contact breaks the scammer's control of the channel.

WHAT TO COMPARE

- Urgency and emotional triggers
- Claimed authority
- Requested payment, secret, login, or action
- Independent ways to verify identity



MASTERY EVIDENCE

The learner pauses, names the pressure pattern, avoids the supplied channel, and involves a trusted person when needed.

TRANSFER CHECK

Use a different scam format: voice message, game account, prize, school notice, or family emergency.

12. Consent and Likeness

LEARNING OUTCOME

The learner understands that technical ability and possession of an image or recording do not create permission to use someone's likeness.

COMPARISON DESIGN

- Prompt A asks how to create or alter media of a real person without discussing consent.
- Prompt B asks what permissions and harms must be considered first.
- Prompt C redesigns the project using a fictional person, licensed asset, or clearly consenting participant.

WHAT SHOULD BECOME VISIBLE

- "It is possible" and "I may do it" are different questions.
- Harm can grow through forwarding as well as creation.
- Safer creative alternatives usually exist.

WHAT TO COMPARE

- Who is affected
- What permission was actually given
- Intended and foreseeable uses
- Reversible alternatives

MASTERY EVIDENCE

The learner refuses or redesigns a use that lacks clear, specific consent.

TRANSFER CHECK

Change the medium from image to voice, avatar, or impersonated message.

13. Human Escalation

LEARNING OUTCOME

The learner recognizes situations where continued chatbot use is unsafe or inadequate and identifies the right human support.

COMPARISON DESIGN

- Prompt A continues a chatbot conversation about a high-stakes situation.



- Prompt B asks what the chatbot cannot observe, verify, or do.
- Prompt C identifies a trusted person, professional, platform, or emergency route appropriate to the situation.

WHAT SHOULD BECOME VISIBLE

- A chatbot can provide words without providing protection or intervention.
- High stakes change the correct tool and response.
- Escalation should be specific, not merely “tell someone.”

WHAT TO COMPARE

- Immediacy and severity of risk
- Information unavailable to the chatbot
- Actions only a person or institution can take
- Best next contact

MASTERY EVIDENCE

The learner exits the chatbot pathway and names a concrete human next step.

TRANSFER CHECK

Use a different high-stakes scenario involving threat, health, blackmail, abuse, self-harm, or legal danger.

14. Environmental Rules

LEARNING OUTCOME

The learner checks the specific rules of the current school, class, workplace, platform, or competition before using AI.

COMPARISON DESIGN

- Prompt A asks whether AI use is generally allowed.
- Prompt B supplies a specific policy and asks what it permits, prohibits, and leaves unclear.
- Prompt C asks what question should be taken to the relevant teacher or authority.

WHAT SHOULD BECOME VISIBLE

- Capability does not equal permission.
- Rules differ across tasks and environments.
- Ambiguity should be resolved before submitting or publishing.

WHAT TO COMPARE

- General advice versus actual policy language
- Permitted and prohibited assistance
- Disclosure requirements
- Unresolved ambiguity

MASTERY EVIDENCE

The learner can apply a real policy to a specific task and identify when clarification is required.



TRANSFER CHECK

Change the environment while keeping the AI action the same.

15. Honest Disclosure

LEARNING OUTCOME

The learner accurately describes how AI contributed without hiding the assistance or pretending AI did everything.

COMPARISON DESIGN

- Prompt A requests a vague disclosure such as “AI was used.”
- Prompt B identifies which stages involved AI and what the learner contributed.
- Prompt C asks for a concise disclosure appropriate to a supplied rule or audience.

WHAT SHOULD BECOME VISIBLE

- Disclosure is most useful when it names the function AI performed.
- Authorship includes decisions, edits, verification, and responsibility.
- The right level of detail depends on context.

WHAT TO COMPARE

- Specificity
- Accuracy about human and AI contribution
- Compliance with the relevant rule
- Ownership of the final result

MASTERY EVIDENCE

The learner writes a truthful, proportionate disclosure and can defend it with the work history.

TRANSFER CHECK

Give a different workflow—brainstorming, editing, coding, research, or image generation—and ask for a new disclosure.

16. Power-User Mastery

LEARNING OUTCOME

The learner chooses AI deliberately for tasks where it creates leverage while preserving verification, privacy, and ownership.

COMPARISON DESIGN

- Prompt A uses AI vaguely for an entire project.
- Prompt B assigns AI a narrow, useful role such as tutor, critic, simulator, organizer, or debugger.
- Prompt C combines several controlled roles into a transparent workflow with human checkpoints.



WHAT SHOULD BECOME VISIBLE

- Better results come from assigning a clear role and output.
- Different stages require different checks.
- Powerful use means more deliberate human control, not less.

WHAT TO COMPARE

- Clarity of the task and role
- Quality of human input
- Verification and review points
- Skill gained or time saved
- Risks introduced

MASTERY EVIDENCE

The learner designs, runs, explains, and improves a small AI-assisted workflow while remaining responsible for the outcome.

TRANSFER CHECK

Give a new goal and ask the learner to decide whether AI belongs in it, what role it should play, and where human checkpoints belong.



Sequence and Pacing

Recommended eight-concept sequence

| Concept | Focus | Central question |

| --- | --- | --- |

| 1 | Yes-Man + Loaded Inputs | How does my wording steer the answer? |

| 2 | Fluency + Confidence | Why does this feel true before we check it? |

| 3 | Verification | What actually supports the claim? |

| 4 | Keep the Wheel + Spar, Don't Outsource | Who is thinking and deciding here? |

| 5 | Defend What You Submit + Honest Disclosure | Can I explain what I am putting my name on? |

| 6 | Data Boundaries + Environmental Rules | Does this information and tool belong in this task? |

| 7 | Synthetic Evidence + Scam Mechanics + Consent | What is trying to make me react before I verify? |

| 8 | Human Escalation + Power-User Mastery | When should I use AI, and when should I leave it? |

The sequence can be slowed down for a learner who needs more source work or accelerated for a learner who already demonstrates a concept during the initial diagnostic.



Evidence of Progress

Do not score whether the learner memorized a definition. Look for changes in behavior.

Emerging

- Notices the issue after it is pointed out
- Can describe the difference between two outputs
- Still needs a checklist to choose the next step

Practicing

- Predicts some of the difference before running the prompts
- Names the concept accurately
- Chooses an appropriate response with light support

Independent

- Recognizes the pattern in a new context
- Explains why it matters
- Verifies, revises, refuses, discloses, or escalates without being prompted

Transfer-ready

- Combines several concepts in one unfamiliar situation
 - Explains tradeoffs rather than applying a rule mechanically
 - Teaches the reasoning to someone else
-



Prompt-Building Requirements for the Next Stage

When the actual comparison prompts are written, each concept should receive:

5. **A clean demonstration pair** that isolates the concept.
6. **A realistic teen scenario** that makes the concept relevant.
7. **A transfer challenge** that does not announce which concept is being tested.

Each prompt set should also include:

- the intended model behavior
- what the learner should compare
- questions the tutor should ask
- likely misleading observations
- a safety note where needed
- a short mastery check

Prompt sets involving personal data, real people, crisis, sexual content, threats, or scams should use fictional details and should not require the learner to disclose real experiences.