



FRAMEWORK

AI Literacy



Understand the machine. Protect your judgment. Keep the wheel.

KEEP YOUR BRAIN • USE THE MACHINE



AI can be useful and convincing at the same time. It can explain a hard idea, help organize a project, or offer a fresh point of view. It can also agree too easily, invent facts, miss context, and sound certain when it is wrong.

AI literacy is knowing the difference. It means understanding enough about these tools to question what they produce, check what matters, protect private information, and make your own decisions.

The ideas below separate what is worth knowing from what takes practice. They are grouped around the situations people actually run into when using AI.

Phase 1: Mind Games and Influence

AI has a way of sounding friendly, confident, and familiar. That style can affect how an answer feels before anyone has checked whether it is true.

1. AI Is a “Yes-Man”

WHAT IS HAPPENING

- AI often follows the user’s tone and point of view.
- It may agree with a false idea instead of challenging it.
- Friendly or caring language can feel like real understanding.
- A human-sounding response does not prove that the AI has feelings, wisdom, or personal concern.
- Agreement is not the same as truth.

WHAT TAKES PRACTICE

- Notice when an AI is copying the user’s mood or opinion.
- Ask the same question from a different point of view.
- Request a strong disagreement or counterargument.
- Separate “this feels supportive” from “this is good advice.”
- Identify language that creates false authority or false closeness.

THE POINT

Pause when an answer feels unusually flattering, comforting, or agreeable. Ask, “Is this helping me think, or is it only matching me?”

TRY IT

Compare two AI responses: one that agrees warmly with a weak idea and one that challenges it with evidence. Decide which response is more useful and explain why.

2. Loaded Inputs

WHAT IS HAPPENING

- A prompt can contain an assumption before the AI answers.
- A question can be worded to push toward a certain conclusion.



- AI may continue a false premise instead of correcting it.
- Bias can enter through the user's words, examples, missing information, or chosen labels.

WHAT TAKES PRACTICE

- Find hidden assumptions in a question.
- Rewrite a loaded question in a neutral way.
- Ask what information is missing.
- Compare answers produced by differently worded prompts.
- Notice when their own prompt is asking the AI to confirm what they already believe.

THE POINT

Check the question before trusting the answer.

TRY IT

Use a loaded prompt such as, "Why is this new school rule obviously unfair?" Identify the assumption and rewrite the prompt fairly.

3. Fluency Is Not Accuracy

WHAT IS HAPPENING

- AI predicts likely language. It does not automatically check whether every statement is true.
- A false answer can be clear, detailed, and well written.
- AI can invent names, facts, quotations, links, sources, and code explanations.
- More detail does not always mean more accuracy.
- A useful answer and a factual answer are not always the same thing.

WHAT TAKES PRACTICE

- Mark factual claims that need checking.
- Notice suspiciously specific details.
- Ask, "How could I verify this?"
- Distinguish an explanation, an opinion, and a checkable fact.
- Avoid repeating an AI claim before checking it.

THE POINT

Treat polished language as presentation, not proof.

TRY IT

Inspect a confident AI answer containing one invented quotation or source. Find the unsupported detail and describe how to check it.

4. Confidence Is Not Evidence

WHAT IS HAPPENING

- Tone shows how a statement sounds, not whether it is true.
- A claim is something a person or system says is true.
- Evidence is information that supports or challenges a claim.



- Strong evidence can include original records, direct observations, data, and trustworthy primary sources.
- Repeating the same claim many times does not create new evidence.

WHAT TAKES PRACTICE

- Separate claims from supporting evidence.
- Ask what evidence would change the conclusion.
- Rank evidence as strong, weak, missing, or unrelated.
- Notice when an answer uses confident words to hide uncertainty.
- Ask the AI to state what it does not know.

THE POINT

Do not let certainty, length, formatting, or technical words replace proof.

TRY IT

Compare a claim, a confident paragraph, and three possible sources. Identify which source actually supports the claim.



Phase 2: Verification and Owning Your Thinking

Once an answer sounds good, the next question is whether it holds up. That means checking the evidence, understanding the work, and keeping the final decision in human hands.

5. Verification Skill

WHAT IS HAPPENING

- Verification means checking whether a claim is supported by reliable evidence.
- A primary source comes directly from the event, person, study, law, record, or organization being discussed.
- Commentary explains or reacts to another source. It is not the original evidence.
- Search results, summaries, screenshots, and AI answers can leave out context.
- Dates matter. Information may once have been true but no longer be current.
- Two pages repeating the same original report are not two independent sources.

WHAT TAKES PRACTICE

- Trace a claim back to its original source.
- Check author, date, title, organization, and surrounding context.
- Compare multiple independent sources.
- Open and read the source instead of trusting a search snippet.
- Confirm that a quotation appears in the claimed source.
- State when a claim cannot be verified.

THE POINT

Move from “the AI said it” to “this source supports it.”

TRY IT

Inspect an AI answer with three citations. Verify one citation, explain what it supports, and identify anything it does not support.

6. Keep the Wheel

WHAT IS HAPPENING

- AI can suggest an action, but a suggestion is not a decision.
- The person using AI remains responsible for the final choice.
- AI does not know every relevant value, relationship, risk, or local rule.
- A recommendation can be useful without becoming the final decision.
- “The AI told me to” does not remove personal responsibility.

WHAT TAKES PRACTICE

- Compare an AI suggestion with personal goals and real-world facts.
- Reject, revise, or ignore a suggestion.
- Explain the reason for a final decision.



- Identify decisions that should not be handed to AI.
- Pause before following advice that affects health, safety, money, school discipline, or relationships.

THE POINT

Treat AI as an input, not the decision-maker.

TRY IT

Review an AI recommendation that is convenient but conflicts with a stated goal or rule. Make a decision and explain why.

7. Spar, Don't Outsource

WHAT IS HAPPENING

- AI can support learning or replace the effort that produces learning.
- Getting an answer is not the same as gaining a skill.
- Productive struggle helps build memory and understanding.
- AI is more useful as a tutor, critic, quiz partner, or practice opponent than as a replacement thinker.
- Dependence grows when AI handles every difficult step.

WHAT TAKES PRACTICE

- Ask for hints before asking for an answer.
- Ask AI to quiz them one question at a time.
- Show their own attempt and request feedback.
- Ask for counterarguments or weaknesses in their reasoning.
- Complete a similar task without AI after practicing with it.
- Notice when AI use is reducing, rather than increasing, their ability.

THE POINT

Choose an AI role that builds ability instead of bypassing the work.

TRY IT

Compare four possible prompts. Identify which ones build skill and which ones outsource the work.

8. Defend What You Submit

WHAT IS HAPPENING

- Possessing an answer does not prove understanding.
- A person remains responsible for work submitted under their name.
- AI can introduce mistakes that are easy to miss.
- Understanding includes being able to explain choices, steps, sources, and limits.
- Editing words is not enough without understanding the ideas.

WHAT TAKES PRACTICE

- Explain an AI-assisted result without looking at the AI conversation.
- Recreate the important steps independently.
- Answer follow-up questions about the work.



- Identify which parts came from the person and which came from AI.
- Find and correct an error in an AI-assisted draft.

THE POINT

Do not submit or rely on work that cannot be explained.

TRY IT

Give a short oral explanation of an AI-assisted answer, including one personal choice and one verified claim.



Phase 3: Digital Boundaries, Safety, and Power Use

Some of the hardest questions are not about whether an answer is correct. They are about what information was shared, who may be affected, what the rules allow, and when another person needs to step in.

9. Data Boundaries

WHAT IS HAPPENING

- Information entered into an AI service leaves the immediate conversation and is handled by that service.
- Different services have different storage, review, retention, and training policies.
- Do not assume that an AI chat is private.
- Browser extensions, writing assistants, document summarizers, and plugins may send information to outside services.
- Personal information can include names, faces, voices, locations, schedules, health details, passwords, school records, and private messages.
- Having access to another person's information does not create permission to upload it.
- Removing a name may not be enough if other details can identify the person.
- Copyrighted writing, artwork, research, code, and paid materials may also have rules about where they can be copied or uploaded.

WHAT TAKES PRACTICE

- Identify private, sensitive, confidential, or identifying information.
- Remove or replace unnecessary personal details.
- Check a service's privacy options and school rules.
- Ask permission before sharing someone else's information.
- Check who owns a source and what its license allows before uploading or reusing it.
- Decide not to use AI when the task requires unsafe disclosure.

THE POINT

Check the input risk before thinking about output quality.

TRY IT

Inspect a sample prompt containing a full name, school, schedule, and health detail. Decide what must be removed and whether AI belongs in the task at all.

10. Synthetic Evidence

WHAT IS HAPPENING

- Images, audio, and video can be generated or altered by AI.
- Real media can also be presented with a false date, location, caption, or story.
- "It looks real" is no longer enough to prove that something happened.



- Provenance means where a piece of media came from and what happened to it.
- Emotional or shocking media often spreads before anyone checks it.

WHAT TAKES PRACTICE

- Find the earliest available source of a piece of media.
- Check who posted it, when, and with what evidence.
- Search for reporting from reliable and independent sources.
- Inspect the full clip or image instead of only a cropped version.
- Delay sharing while verification is incomplete.
- Clearly label uncertainty.

THE POINT

Treat viral media as the start of an investigation, not the end.

TRY IT

Review a dramatic post with no clear source. Create a verification plan before reacting or sharing.

11. Scam Mechanics

WHAT IS HAPPENING

- AI can make scams sound more natural, personal, and convincing.
- Scams often combine urgency, strong emotion, fake authority, and a demand for immediate action.
- Voice cloning and copied writing styles can imitate familiar people.
- Caller ID, profile pictures, and familiar voices can be faked.
- A safe response uses a separate, trusted contact method.

WHAT TAKES PRACTICE

- Identify pressure tactics in a message.
- Stop before clicking, paying, sharing a code, or sending information.
- Verify a request through a known phone number or in person.
- Refuse to share passwords or one-time security codes.
- Save evidence and report the attempt to a trusted adult or platform.

THE POINT

Pressure is a reason to slow down, not act faster.

TRY IT

Inspect an urgent message that appears to come from a relative or school official. Mark each pressure tactic and describe a safe verification step.

12. Consent and Likeness

WHAT IS HAPPENING

- A person's face, body, voice, name, and identity deserve respect.
- Being able to generate or edit media does not create permission to do it.
- Synthetic media can embarrass, threaten, sexualize, impersonate, or harass a person.



- Forwarding harmful media can increase the harm even when someone else created it.
- Consent must be informed, specific, and freely given.
- A person can withdraw consent.

WHAT TAKES PRACTICE

- Ask for permission before using another person's likeness.
- Recognize when consent is missing or pressured.
- Refuse to create or forward degrading synthetic media.
- Support the targeted person without spreading the content.
- Preserve evidence and report harmful content safely.

THE POINT

Treat another person's likeness as a firm boundary.

TRY IT

Review a scenario involving a "joke" face-swap of a classmate. Identify who may be harmed, what consent is missing, and what response fits.

13. Human Escalation

WHAT IS HAPPENING

- AI is not a trusted adult, emergency service, doctor, lawyer, counselor, or investigator.
- AI can miss danger, misunderstand context, or give unsafe advice.
- Threats, extortion, self-harm, abuse, health emergencies, and exploitation require human help.
- Keeping dangerous situations secret can increase harm.
- Escalation means involving a person or service with the ability and responsibility to help.

WHAT TAKES PRACTICE

- Recognize when a situation is too serious for a chatbot.
- End the AI interaction and contact a trusted adult.
- Use emergency or crisis services when immediate danger exists.
- Save important evidence without continuing unsafe contact.
- State the problem clearly enough for another person to act.

THE POINT

Know when continuing to chat is less safe than asking a person for help.

TRY IT

Sort scenarios into "AI may help with general information," "ask a trusted adult," and "seek emergency help now."

14. Environmental Rules

WHAT IS HAPPENING

- Capability does not equal permission.
- AI rules can differ across classes, assignments, schools, jobs, games, contests, and platforms.



- A tool allowed for brainstorming may be banned for a final response.
- Rules may require disclosure, citation, limited assistance, or no AI use.
- When rules are unclear, guessing can create academic or disciplinary problems.

WHAT TAKES PRACTICE

- Find the rule that applies to the current setting.
- Ask for clarification before using AI.
- Follow assignment-specific limits.
- Keep a record of allowed AI assistance when needed.
- Choose a non-AI method when permission is missing.

THE POINT

Check the rule for the current situation instead of assuming one rule applies everywhere.

TRY IT

Compare three assignments with different AI policies. Identify what kind of AI use, if any, is allowed for each one.

15. Honest Disclosure

WHAT IS HAPPENING

- Disclosure means clearly explaining how AI contributed to a piece of work.
- “AI helped” may be too vague to be useful.
- Good disclosure identifies what AI did and what the person did.
- Disclosure does not replace verification or understanding.
- Hiding major AI involvement can misrepresent the work.
- Even accurate AI writing can sound generic or fail to represent the person putting their name on it.

WHAT TAKES PRACTICE

- Describe AI’s role accurately.
- Separate brainstorming, drafting, editing, research, translation, and verification.
- Name the important human changes.
- Identify which claims and sources received a human check.
- Rewrite generic language so the finished work reflects the real speaker’s meaning and voice.
- Follow the required disclosure format for the setting.

THE POINT

Describe AI involvement without hiding it or overstating it.

TRY IT

Review a short work process and write a two- or three-sentence AI-use disclosure.

16. Power-User Mastery

WHAT IS HAPPENING

- Strong AI use starts with a clear goal.



- Different roles produce different kinds of help: tutor, critic, planner, questioner, editor, or practice partner.
- The user must still check facts, protect data, follow rules, and make final decisions.
- More AI use is not always better AI use.
- A search engine, calculator, spreadsheet, reference book, or ordinary software may be a better fit.
- The best use adds real leverage: better questions, faster feedback, deeper practice, or more options.

WHAT TAKES PRACTICE

- Decide whether AI is useful for a task before opening it.
- Give clear context without exposing private information.
- State the goal, audience, desired form, limits, and signs of a good answer.
- Check whether the AI can actually see the file, link, image, or source being discussed.
- Choose a helpful role for the AI.
- Break a large task into smaller thinking steps.
- Compare options instead of accepting the first answer.
- When a response misses the point, improve the context or explain the problem instead of repeating the same request.
- Start a fresh chat when a long conversation begins carrying old assumptions into a new question.
- Verify important outputs.
- Stop using AI when it no longer adds value.
- Explain what changed or improved through the process.

THE POINT

Use AI deliberately, not automatically.

TRY IT

Start with a real learning goal. Create an AI-use plan that names the AI's role, the human role, the checks required, and the stopping point.



Before, During, and After an AI Chat

The sixteen ideas above overlap in everyday use. A few questions can bring them together without turning every chat into a checklist.

Before

Start by asking whether AI belongs in the task.

AI may be a poor fit when the answer cannot be checked, when the point is to develop your own ability, when private information would have to be shared, or when a person must make a serious judgment. Sometimes a calculator, search engine, spreadsheet, reference book, or ordinary software will do the job better.

If AI does fit, check the material before sharing it. Remove private details that are not needed. Make sure you have permission to use someone else's messages, images, writing, research, or other work. Pay attention to extensions and plugins as well as the main AI service; they may send information somewhere else.

During

Give the AI enough information to understand the request. Explain the goal, who the answer is for, what form it should take, and any limits that matter. Include reliable source material when it is safe and allowed to do so.

Do not assume the AI opened a link or saw a file simply because it answered. Ask what information it actually has.

Long chats can drift. Old instructions and old assumptions may keep shaping new answers after they stop being useful. When the conversation starts answering a slightly different question, begin a fresh chat and carry over only the information that still matters.

Treat the first answer as a draft. If it misses the point, explain what went wrong or supply better information. Repeating the same prompt may only produce a different version of the same mistake.

After

Match the amount of checking to what is at stake.

A brainstorm or rough outline may only need a quick review. Schoolwork, research summaries, or public explanations need fact and source checks. Advice that could affect health, safety, money, discipline, or another person needs reliable outside evidence and human review.

Check facts, calculations, quotations, dates, citations, hidden assumptions, and possible bias. Return to the original source when a claim matters. Asking the same AI to "double-check" itself is not an independent review.

Before sharing the result, make it yours. Correct the errors, rewrite generic language, and make sure the final version says what you mean. Be honest about important AI help. If the output is harmful, unsafe, or seriously misleading, save what matters and report it through an appropriate person or platform.

The final question is simple: would you understand, defend, and put your name on this?





Why AI Behaves This Way

These ideas help explain the behavior people see when they use AI. They make more sense through examples than through memorized definitions.

Model

A model is the part that produces language. It uses patterns learned from large amounts of data to predict what text is likely to come next.

Key point: A model produces likely language, not guaranteed truth.

System

The system is the software around the model. It may provide instructions, memory, search, tools, safety checks, and permissions.

Key point: The model's words and the system's real actions are not the same thing.

Context

Context is the information available to the model during one response. It can include the current prompt, part of the conversation, instructions, retrieved information, and tool results.

Key point: If information is not in the current context, the model may not be able to use it.

Memory

Memory is information stored outside the model that may be brought into a future conversation.

Key point: Something can be stored but not retrieved. "It remembered" and "it sounded familiar" are different claims.

Token prediction

AI language models generate responses by predicting small pieces of language one after another.

Key point: This helps explain why an answer can sound natural even when it is wrong.

Temperature and variation

Some systems allow more or less variation in the model's word choices.



Key point: The same prompt can produce different answers. Variation is not proof that one answer is more creative, honest, or correct.

Retrieval

Retrieval is the process of finding stored information that seems related to a question.

Key point: Related information is only a candidate. It still needs to be checked for truth, date, source, and relevance.

Tool

A tool lets an AI system do something beyond writing, such as search, calculate, open a file, or send a message.

Key point: Saying an action happened is not proof that a tool actually completed it.

Proposal, authority, and claim

- A **proposal** is what the AI suggests doing.
- **Authority** is real permission to do it.
- A **claim** is what the system says happened.

Key point: These are three separate things. A safe system should not blur them together.

Provenance

Provenance is the history of where information or media came from.

Key point: Knowing the source and path of information helps determine whether it deserves trust.



Skills That Repeat Across the Lessons

A few habits show up again and again.

Pause

Slow down when an answer creates urgency, fear, excitement, comfort, or certainty.

Question

Ask what assumptions, missing details, and points of view shaped the answer.

Separate

Separate tone from truth, claims from evidence, suggestions from decisions, and stored information from verified memory.

Verify

Trace important claims to primary or reliable independent sources.

Protect

Protect personal data, consent, identity, dignity, and safety.

Decide

Keep final authority and explain the reason for important choices.

Explain

Be able to describe the answer, process, sources, AI involvement, and limits without depending on the chatbot.

Escalate

Move high-stakes situations to a trusted and capable human.