Teacher Edition Lesson Pack

CivicOS Labs

Lesson 1: What AI Is and Is Not

Teacher Edition

Lesson at a Glance

Lesson length: 45 minutes core instruction + 20 minutes student independent work (in-class or homework)

Learning objectives. By the end of this lesson, students will be able to: 1. Explain in their own words what a Large Language Model (LLM) is and how it works at a basic level 2. Identify three things AI is good at and three things AI is not good at 3. Define and apply the term “hallucination” to an AI context 4. Articulate why verification of AI answers matters 5. Produce an AI Concept Map that demonstrates conceptual understanding

Florida standards addressed: - **FL B.E.S.T. CS (grades 6–8):** Emerging Technologies — AI benchmarks (SC.6.ET.2.1, SC.6.ET.2.2, SC.7.ET.2.1, SC.8.ET.2.1, SC.8.ET.2.2, primary); **Programming and Software Engineering — model limits (SC.8.PE.3.3, primary substance fit for “what AI is not good at”)**; Communication and Collaboration (SC.7.CC.2.2); Computing Components (SC.7.CO.2.2) - **Civics:** SS.7.CG.2.9 (analyzing media for bias, symbolism, and propaganda) — *introduced* through hallucination + verification framing; full development in Lessons 2 and 3 - **B.E.S.T. ELA (grades 6–8):** Reading Informational Text — Structure (ELA.7.R.2.1), Vocabulary Acquisition (ELA.7.V.1.1, ELA.7.V.1.3), Communication — Oral (ELA.7.C.2.1), K-12 ELA Expectations — Collaborative Discussion (ELA.K12.EE.4.1)

(See the Standards Alignment Matrix for full alignment detail.)

Materials needed: - Teacher device with browser access for AI demonstration - Class display (projector, smart board, or large monitor) - Student copies of the Lesson 1 Student Edition (printed or digital) - Concept map worksheet or blank paper for AI Concept Map activity - Optional: students’ own devices for follow-up activities

Google Classroom quick access: Google Classroom Upload Pack — setup guide, copy/paste assignment posts, upload manifest, and Bellringer / Warm-Up pacing options.

Prerequisites: None. This is the entry-point lesson.

Vocabulary introduced: AI (Artificial Intelligence), LLM (Large Language Model), Model, Prompt, Hallucinate

Portfolio artifact produced: AI Concept Map (Lesson 1 contribution to the curriculum portfolio)

Assessment: Self-Check questions at end of Student Edition + AI Concept Map evaluated with the rubric in this Teacher Edition

---

Background for the Teacher

This is the foundational lesson. Everything in subsequent lessons builds on what students grasp here. Two pieces are most important to get across:

First, the working model of how an LLM works. Students should leave understanding that an LLM is, at its base, a "predict the next word" engine — and that this seemingly simple mechanism, applied at huge scale, is what makes AI capable of doing so much. This is not just a fact to memorize; it is the conceptual handle that makes everything else in the curriculum make sense. When students later encounter AI hallucination, AI's inability to do math, AI's lack of current information — all of those make sense as consequences of "predict the next word" rather than as separate mysteries.

Second, hallucination as a structural property. This is the most important thing students will learn in the entire curriculum. AI does not occasionally make small mistakes. It produces confidently wrong answers as a routine output of how it works. Students who internalize this — who develop the verification habit early — use AI well throughout their education. Students who do not internalize it tend to get burned at some point, often badly. The lesson treats hallucination seriously and gives students the language to recognize and respond to it.

A common pitfall when teaching this lesson is being either too dismissive ("AI is just statistics; it doesn't really do anything") or too reverent ("AI is amazing and can do almost anything"). Both framings produce worse outcomes than curriculum's framing, which is something like: AI is a powerful tool with specific strengths and specific structural limitations; using it well requires understanding both.

You do not need deep AI expertise to teach this lesson well. You need the working model in this Background section and the lesson plan that follows. Your job is to deliver the conceptual framework and facilitate students' engagement with it; you do not need to be able to answer every advanced AI question that comes up.

If a student asks something you genuinely do not know: model what you want them to do. "I don't know — let me check." Look it up with them. The verification habit you are teaching applies to you too.

Lesson Procedure

Total time: 45 minutes core instruction. Independent work afterward: 20 minutes.

Timing note for first delivery: Times are targets, not requirements. First-time delivery often runs 5–10 minutes long because students ask clarifying questions about what counts as AI. If pacing gets tight, use Bellringer Option A, shorten share-outs, and keep the Concept Map as independent completion work.

Bellringer / Warm-Up Options (choose one)

Use one option at the start of class. If pacing is tight, use Option A and treat it as the first three minutes of the Opening rather than an added activity. Post or read the exact prompt; do not assume students know which tools count as AI.

Option A — Fast Start (3 minutes): Post: "Name one AI tool or AI-powered feature you have used or seen. Examples: ChatGPT, Siri, Alexa, Gemini, Copilot, phone autocorrect, YouTube recommendations, translation apps, or game bots. What do you think it is doing behind the screen?" Students answer in one sentence. Ask

for one quick share-out.

Option B — Standard Warm-Up (5 minutes): Post: "1. Name one AI tool or AI-powered feature you have used or seen. Examples: ChatGPT, Siri, Alexa, Gemini, Copilot, phone autocorrect, YouTube recommendations, translation apps, or game bots. 2. What do you think that tool is doing behind the screen to produce its answer, suggestion, or action? 3. What is one thing you would want to know before trusting its answer or recommendation?" Students answer, then pair-share one answer. Use one response to transition into: "Can someone explain what AI actually is?"

Option C — Extended Launch (6-7 minutes): Post: "1. Name one AI tool or AI-powered feature you have used or seen. 2. What do you think that tool is doing behind the screen? 3. What is one thing you would want to know before trusting its answer or recommendation? 4. Which matters more when using AI: knowing what it can do well, or knowing where it fails? Explain your choice in 2-3 sentences." Let two students take different positions, then frame the lesson as learning both.

Opening (5 minutes)

Begin with a question to the class: "How many of you have used AI before — ChatGPT, Claude, Siri, Alexa, anything similar?" Get a show of hands. Most will have.

Then ask: "Can someone explain what AI actually is?" Let two or three students try. Most explanations will be vague or incomplete; that is the point. Acknowledge: "This is normal. Almost everyone uses AI without being able to explain what it actually is. We're going to fix that today."

State the lesson goal: "By the end of class, you should be able to explain what AI is to someone who has never thought about it — a parent, a younger sibling, a friend. You should also know where AI is reliable and where it isn't, which matters for how you use it."

Direct Instruction (15 minutes)

Have students follow along in their Student Edition (printed or digital). Walk through the reading section by section. Read the first paragraph or two aloud yourself; this models the expected attention level and gives students who read more slowly time to catch up.

The key sections to emphasize:

"The Kind of AI This Lesson Is About" — make sure students understand that "AI" in this curriculum specifically means LLMs. Other kinds of AI exist (image recognition, speech recognition, game-playing AI), but this curriculum is about the kind they are most likely to encounter for school work.

"What an LLM does" — this is the conceptual core. The "predict the next word" framing is what students must internalize. Use the analogy in the reading (someone who has read every book ever) and add a personal example. Try: "Imagine I asked you 'The capital of France is…' What word comes next? Right — Paris. You did the same kind of pattern-matching the AI does. The AI is just doing it on a much bigger scale, with way more text in its memory."

"What AI is good at" and "What AI is not good at" — read these together. Pause after the four limitations and check for understanding: "Why might base LLMs not have current information?" "Why is base AI not reliable for math?" Most students will already understand intuitively that base AI does not know last week's

news; the math limitation often surprises them. Take the time to explain — base LLMs generate digits the same way they generate words, by pattern, not by actually computing. Note that some modern systems connect to tools (web search, calculators, code execution) that improve these limits substantially; the lesson framing acknowledges this rather than asserting absolute incapability.

Hallucination — this is the most important paragraph in the entire lesson. Read it slowly. Pause after. Ask: "What's the difference between AI being wrong sometimes and AI hallucinating?" The answer: AI usually does not *know* it is wrong when it hallucinates. It produces wrong information with the same confident voice as right information. This is what makes hallucination dangerous; the AI cannot reliably warn you. Modern AI systems can reduce hallucination through grounding techniques, but they cannot eliminate it — that is the framing the curriculum uses, and it is the responsible level of claim.

Activity 1: AI Demonstration (10 minutes)

Now demonstrate AI in real time on the class display. Choose a topic that interests your students — sports, history, science, current pop culture (within the AI's knowledge cutoff). Type a prompt and let the class watch the response appear.

Suggested prompts (use one): - "Explain photosynthesis to a 7th grader." - "Tell me three interesting things about the Roman Republic." - "Explain how the body fights off a virus."

As the response generates, point out: - "Notice how it streams the words one at a time? That's the AI predicting word after word, in real time." - "Notice how confident it sounds? That's how AI always sounds — confident, whether it's right or wrong."

When the response is complete, ask: - "Did anything in there sound surprising or new to you?" - "If we wanted to check whether the AI got everything right, what would we do?"

The goal is to make hallucination concrete. Even if the AI's response was actually accurate, the question "how would we check?" sets up the verification habit that becomes the focus of Lesson 2.

If you cannot do live AI demonstration, use the prepared example AI outputs from the No-Install Adaptation Notes section below.

Activity 2: AI Concept Map (15 minutes — most of the independent work time)

Direct students to the Activity 2 instructions in their Student Edition. Walk through the structure: - Center: AI or LLM - Four branches: what AI is good at, what AI is not good at, how AI works (basic), why we verify - Sub-branches with specific examples - One thing they found surprising

Give students 12–15 minutes to work. Circulate and prompt: "Can you add a specific example?" "What's a more concrete way to say that?" "What's missing from this branch?"

The Concept Map is the portfolio artifact for Lesson 1. It does not need to be polished or visually impressive; it needs to demonstrate that the student understands the lesson's conceptual content.

Closing (5 minutes)

Bring the class back together. Have two or three students share one thing from their concept maps — what they put on a particular branch, what surprised them.

Preview Lesson 2: "Tomorrow we're going to take this idea about hallucination seriously. We're going to take an actual AI answer and check it against real sources, then document what we find. The goal is to start building the verification habit that will protect you from the most common way students get into trouble with AI."

Direct students to the Reflection Questions in their Student Edition (5 minutes after class or as homework) and the Self-Check (also after class).

Student Reading — Reproduced with Teacher Annotations

The full Student Edition reading appears below with margin annotations for educators. The annotations flag important points, predictable confusion, and discussion opportunities.

> **You have probably used AI before.** [TEACHER NOTE: Most students will recognize at least one of the AI examples that follow. Ask early in this paragraph for a show of hands; gets students engaged.] **Maybe you asked ChatGPT to help with homework. Maybe you talked to Siri or Alexa. Maybe you got a movie recommendation from Netflix that figured out what you liked. All of those use artificial intelligence.**

> **But here is something interesting. Most people who use AI every day cannot actually explain what it is.** [TEACHER NOTE: This framing is intentional. Students often think their parents or older siblings "understand AI." Most adults don't either. The lesson positions students as gaining understanding most adults lack — a small motivation boost.]

> **When most people say "AI" today, they are usually talking about a specific kind of program called a Large Language Model — an LLM.** [TEACHER NOTE: First introduction of the LLM term. Make sure students hear "Large Language Model" once. Use the abbreviation LLM consistently after that.]

> **Imagine reading every book ever written. Plus most of the internet. Plus millions of articles and conversations.** [TEACHER NOTE: This is the analogy that does most of the conceptual work. Spend time here. Have a student estimate how long it would take a human to read every book ever written. The answer is "many lifetimes." This is what makes the scale meaningful.]

That is the whole job. The AI predicts what word should come next. [TEACHER NOTE: The most important sentence in the lesson. Pause after reading it aloud. Some students will be surprised; some will be skeptical. The skepticism is healthy — it leads naturally into the next section.]

You might be thinking: that does not sound powerful enough to do everything AI can do. [TEACHER NOTE: This paragraph addresses the natural skepticism. The "predict what comes after 'Here is a summary:'" example is the key. Students often have an "oh, I see" moment here.]

First, base LLMs do not have current information unless they're connected to other tools. [TEACHER NOTE: This limitation is intuitive for most students. The "unless they're connected to other tools" clause matters — students may have used ChatGPT or similar with web search enabled and gotten current information; the lesson honors that without overstating the base LLM's capabilities. If a student says "but ChatGPT told me about something that happened yesterday," that's a tool-augmented system at work; the underlying model still doesn't know last week's news on its own.]

Second, AI is not reliable for math unless it's using a calculator or code tool. [TEACHER NOTE: This one usually surprises students. They've seen AI write equations and assume it's actually computing. The key clarification: base LLMs generate digits as patterns, not as computed values. Some modern AI systems do call out to calculators or code execution tools — when that happens, math reliability improves substantially. The lesson's framing is honest about both: base AI is unreliable for math; tool-augmented AI can be reliable. This is also a natural cross-curricular tie to math: "your calculator does math; a base AI predicts what math should look like."]

Third — and this is the most important one — AI hallucinates. [TEACHER NOTE: THE most important paragraph in the lesson. Read slowly. Pause. The structural point — that hallucination is a persistent risk that can be reduced but not eliminated — must land. Students who internalize this become careful AI users; students who miss this get burned later. The lesson's framing acknowledges that modern systems reduce hallucination through grounding, retrieval augmentation, and tool use; it does not assert that engineering will never improve. What it asserts is that hallucination will not disappear and treating AI as if it never hallucinates is risky.]

Fourth, AI only knows what you tell it. [TEACHER NOTE: This sets up the prompting work in Lesson 4 of the Foundation Edition (and in lessons throughout this curriculum). For now, students just need to know that vague questions get vague answers.]

> **The single most important habit when using AI is verification.** [TEACHER NOTE: This is the bridge to Lesson 2, where verification becomes a practiced skill. Set this up with intent: “We’re going to actually practice this tomorrow.”]

> **People argue about whether AI is really thinking or really understanding.** [TEACHER NOTE: This is the philosophical question students will probably bring up if you don’t address it. The reading addresses it briefly. If a student wants to go deeper, the Extension Options below include this discussion. For most classes, the brief treatment in the reading is enough.]

Vocabulary Teaching Strategies

The five vocabulary terms (AI, LLM, Model, Prompt, Hallucinate) appear throughout the curriculum. Investing in teaching them well in Lesson 1 pays off across all subsequent lessons.

Quick check at start of class: Ask the class which of the five terms they have heard before. Most will have heard “AI” and “ChatGPT” but few will have heard “LLM,” “prompt” (in the AI sense), or “hallucinate” (in the AI sense).

Pre-teach “hallucinate”: Before the reading, ask students what “hallucinate” means in everyday English. The usual answer: seeing or hearing things that aren’t really there. Then preview: “We’re going to learn that AI does something kind of like that. AI makes up information that isn’t there but acts like it is. We use the same word.”

Cognate flagging for ELL students: “Modelo” (model) and “prompt” (less direct cognate but learnable) help Spanish-speaking ELL students. “Inteligencia Artificial” maps directly to AI. Reinforce these orally for ELL students before the reading.

Vocabulary check during the lesson: After the reading, before Activity 1, do a quick verbal check. “What does LLM stand for? What’s a model? What’s a prompt? What does hallucinate mean in our AI context?” Sixty seconds. This consolidates before students move to active work.

Differentiation Notes

The Differentiation Notes for Lesson 1 cover the most common needs. Adapt to your specific students.

Below grade level readers

The Student Reading is calibrated to ~Lexile 1100L. Students reading meaningfully below grade level may need:

- **Read-aloud:** read the reading aloud while students follow along. Slows the pace but ensures comprehension.
- **Vocabulary pre-teaching:** introduce all five vocabulary words before the reading begins, with student-friendly definitions on the board
- **Chunked reading:** stop after each major section, ask one comprehension question, then continue
- **Visual organizer:** for the AI Concept Map activity, provide a pre-formatted template with the four branches already drawn; students just fill in the sub-branches

Above grade level readers

For students reading well above grade level:

- **Independent reading first:** let them read silently while the rest of the class reads with the teacher; then pull them into the discussion
- **Extension reading:** Foundation Edition (*The Open Source Student* v1.4) Chapter 1 is the deeper version of this lesson; share it with students who want to go further
- **Mentor role:** ask them to help explain concepts to peers during the Concept Map activity

English Language Learners (ELL)

For ELL students at intermediate proficiency:

- **Pre-teach vocabulary** orally before the reading begins
- **Provide a Spanish-cognate flag list:** AI = Inteligencia Artificial, model = modelo, etc.
- **Pair ELL students with bilingual peers** during discussion when possible
- **Offer the Concept Map in graphic form** rather than text-heavy

For newcomer ELL students: - **Reduce the reading load** to the section headers and the most important paragraph (the hallucination paragraph) - **Allow the Concept Map to be drawn or annotated in the student's strongest language**, with key terms in English

IEP / 504 — ADHD or Executive Function Support

- Break the 45-minute lesson into shorter visible segments with timers
- Use the lesson's natural sections (Opening, Reading, Activity 1, Activity 2, Closing) as movement-break opportunities
- Provide the Concept Map template pre-formatted to reduce executive function load
- Use a visual progress tracker on the board so students see where they are in the lesson

IEP / 504 — Reading or Language Processing Support

- Provide audio of the Student Reading (built-in TTS or your own recording)
- Allow student to demonstrate understanding through speech rather than written response on the Concept Map
- Provide larger font printable version

IEP / 504 — Autism Spectrum

- Preview the lesson structure visually before starting (the lesson follows the predictable lesson structure: Opening / Reading / Activity 1 / Activity 2 / Closing)
- For the demonstration in Activity 1, give an explicit advance signal: "I'm about to type into the AI; the response will appear on the screen; we'll discuss after"
- The AI subject matter itself may be an area of strong interest; channel that interest into the Extension Options if needed

Anxiety or Emotional Regulation

- Self-Check questions at the end of the Student Edition can be done privately; students do not need to share answers
- The Concept Map is graded on thinking, not artistic quality; reassure students about this directly
- Discussion participation is invited but not required; students can write reflections instead

Gifted students

- Extension Option 1 (the philosophical question) is a good fit
- Extension Option 2 (designing a prompt to reveal hallucination) suits students who want to engage technically
- Cross-Curricular Connection to writing/literature (how is AI's pattern-matching different from how human writers use patterns?) suits literature-strong students

Limited Technology Access (no personal device)

- The whole lesson works on the teacher's single device with class projection; no student device is needed
- Concept Map can be on paper instead of digital
- Self-Check can be answered on paper

Assessment Rubric: AI Concept Map

The AI Concept Map is the portfolio artifact for Lesson 1. Evaluate based on conceptual understanding, not artistic merit.

Criterion	Developing	Proficient	Exemplary
Center concept (AI/LLM) clearly identified	AI mentioned but not centered	AI/LLM clearly central	AI/LLM central with brief definition or analogy
Four required branches present	Two or three branches	All four branches present	All four branches present, clearly labeled, well-organized
Specific examples on each branch	One or two examples; mostly general	At least three specific examples per branch	Multiple specific examples; some examples connect to student's life
"How AI works" branch shows real understanding	Vague reference (e.g., "computers do it")	"Predicts next word" or equivalent	"Predicts next word" framing with the scale element ("on huge amounts of text")

Criterion	Developing	Proficient	Exemplary
“Why we verify” branch shows understanding of hallucination	Verification mentioned but not connected to hallucination	Connection made: AI can be wrong, verification protects against it	Hallucination explained as structural; verification framed as a habit, not a one-time check
One personal “surprised by” element	Missing or generic	One specific thing the student found surprising	Surprised-by element shows real engagement with the concept
Overall coherence	Map is hard to follow	Map is organized; reader can trace the logic	Map shows the student’s own thinking, not just transcription of the lesson

Sample Proficient artifact (for educator reference):

A concept map with “LLM (Large Language Model)” centered, defined briefly as “AI that works with language.” Four branches: “Good at” (with sub-branches: explaining concepts, summarizing, brainstorming, helping with writing); “Not good at” (with sub-branches: current news — does not know, math — pretends but is wrong, made-up facts — hallucination); “How it works” (with sub-branches: trained on huge text, predicts next word, generates one word at a time); “Why verify” (with sub-branches: AI can sound confident but be wrong, checking against real sources catches errors, this is a habit not a one-time thing). Surprised by: “I did not know AI was actually predicting the next word — I thought it was ‘thinking’ more like a person.”

Sample Exemplary artifact: the same content but with cross-connections drawn (e.g., a line from “predicts next word” to “makes up facts” with a note: “this is why hallucination happens”), and a personal application (“I’m going to verify the science homework AI helps with from now on”).

Answer Key / Scoring Notes

Use these notes when reviewing the warm-up, discussion responses, Concept Map, reflection, and Self-Check. Do not require student wording to match exactly; require the core idea.

Bellringer / Warm-Up expected responses: - Acceptable AI examples include ChatGPT, Siri, Alexa, Gemini, Copilot, phone autocorrect, recommendation systems, translation apps, grammar suggestions, game bots, or other AI-powered features. - A proficient response says the tool is using patterns in data to produce an answer, suggestion, recommendation, or action. - A developing response may say “it searches” or “it knows things.” Use that as a teaching moment: search engines retrieve pages; LLMs generate language from patterns unless connected to retrieval tools.

Activity 1 demonstration expected observations: - Students should notice that the AI response appears confident and fluent. - Stronger responses also notice that confidence is not the same as accuracy. - If students ask how to check the response, expected verification sources include a textbook, teacher-approved source, official source, reputable educational website, or expert source relevant to the topic.

Concept Map required content: - Center: AI or LLM. - Four branches: what AI is good at; what AI is not good at; how AI works; why we verify. - “How AI works” branch should include “predicts next word” or equivalent. - “Not good at” branch should include at least two of: current information without tools, exact math without calculator/code tools, fabricated quotes/facts, hallucination, privacy-sensitive judgment, or unverified claims. - “Why verify” branch should connect verification to hallucination, confident wrong answers, or protecting school/civic work from false information.

Self-Check answer key: 1. **LLM** stands for **Large Language Model**. 2. Strong answers name one language task AI can often support, such as explaining, summarizing, brainstorming, drafting, translating, or reorganizing text; and one limitation, such as current information without tools, exact math without tools, privacy judgment, or hallucination. 3. AI "hallucinates" when it produces information that sounds plausible but is false, invented, or unsupported, usually without reliably warning the user. 4. Verification matters because AI can sound confident while being wrong; real sources help confirm, correct, or reject the claim. 5. Do not trust a famous-person quote just because AI gives it. Verify it against a reliable quote archive, primary source, official collection, reputable biography, or another source that identifies where the quote originally appeared.

Reflection scoring notes: - Grade reflections for specificity and honesty, not for agreeing with a particular view of AI. - A proficient reflection names a concrete change in understanding and a specific responsible or irresponsible use. - An exemplary reflection connects the lesson to future behavior, for example: "I can use AI to brainstorm, but I should verify facts and not paste private information."

Extension Options

For students who finish the Concept Map quickly or want to go deeper:

Extension 1: The Philosophical Question. Have the student write a short response (1 page) to the question "Is AI really thinking?" Should reference the patterns/understanding distinction from the reading. Particularly suitable for verbally strong students or those interested in philosophy.

Extension 2: Design a Hallucination Test. Have the student design a prompt that they think might cause AI to hallucinate. Predict what the AI will get wrong, then test it (with teacher demonstration or browser-safe AI). Document what happened. This is a hands-on entry point to Lesson 2's verification work.

Extension 3: Compare AI Tools. If you have access to two different AI services (e.g., ChatGPT and Claude, or two different open-source models), have the student ask the same question of both and compare the responses. What's the same? What's different? What does this tell us about AI?

Extension 4: Foundation Edition Reading. Direct interested students to *The Open Source Student* (Foundation Edition) Chapter 1 for the longer, deeper treatment of the same content. Suitable for students reading at high school level or above.

No-Install Adaptation Notes

The default Lesson 1 design assumes you can demonstrate AI live during Activity 1 using a browser-based commercial AI service. If your setting does not permit live AI access, here is the no-install version of Activity 1.

Pre-prepared example AI outputs (text descriptions in this Teacher Edition; project as text or convert to screenshots if useful):

For the demonstration moment in Activity 1, use these prepared examples instead of live AI generation. (In the published curriculum, screenshots will be embedded; for the Pilot Kit Markdown source, descriptions are provided that can be turned into screenshots or projected text.)

Example 1: AI explains photosynthesis (typical good response)

Prompt: “Explain photosynthesis to a 7th grader.”

Response (typical AI output, summarized for your demonstration): “Photosynthesis is how plants make their own food using sunlight. Plants take in carbon dioxide from the air through tiny holes in their leaves called stomata. They take in water from the soil through their roots. Using energy from sunlight (captured by chlorophyll, the green pigment in leaves), they combine water and carbon dioxide to make sugar (glucose) for food. They release oxygen as a waste product, which is what we breathe. The simple equation is: water + carbon dioxide + sunlight → sugar + oxygen.”

Discussion prompts: - “Did anything in here surprise you?” - “If we wanted to check whether the AI got the facts right, where would we look?” (Expected answer: textbook, science website, teacher.)

Example 2: AI hallucinates a quote (illustration of the hallucination problem)

Prompt: “Give me a quote by Albert Einstein about education.”

Response (commonly hallucinated AI output — this is a real type of hallucination): “Einstein said, ‘Education is what remains after one has forgotten what one has learned in school.’”

Reality: this quote is widely misattributed to Einstein. There is no reliable source confirming he said it. The quote actually comes from a different source (often attributed to B.F. Skinner, though even that is uncertain).

Discussion prompts: - “If I needed this quote for an essay, what should I do before using it?” - “Notice how confident the AI was. Did it warn us that this might be wrong?” (Answer: no.) - “This is what we mean by hallucination.”

This is one of the most powerful demonstrations available and works without any live AI access — it shows hallucination concretely with a real example.

Common Misconceptions

Students often arrive at this lesson with one or more of the following misconceptions. Address them when they surface; do not assume the lesson alone catches all of them.

“AI knows everything.” The lesson directly addresses this with the four limitations. Reinforce: AI is bounded by its training data; it does not have current information; it cannot reliably do math; it can be wrong about anything.

“AI is just like a search engine.” AI does not search for information when you ask; it generates a response from patterns. A search engine looks up real pages; AI predicts what an answer might look like. This is why AI can confidently produce wrong answers — it isn’t actually retrieving anything.

“AI is going to take my job / my parent’s job.” This concern surfaces sometimes, especially with older students. The honest answer: AI changes what jobs look like; some jobs change a lot; some go away; new ones emerge. The students best positioned for what comes next are the ones who understand AI well enough to use it as a tool, not those who fear it or those who treat it as magic. This curriculum is designed to produce that first kind of student.

“AI is alive / AI is conscious.” Some students treat AI as having feelings or awareness. The reading addresses the philosophical question briefly. The practical answer: AI is software running on a computer. It does not have feelings. Treating it as if it does does not help the student use it well.

“My parents told me not to use AI.” Some parents are skeptical; some are forbidding. Validate the parent’s concern (they’re trying to protect the student) without endorsing avoidance. The curriculum’s stance: use AI well, with verification, ethics, and judgment. That stance is what most parents actually want, even if “don’t use AI” was their first reaction. The Parent / Guardian Letter in the Implementation Readiness Packet is designed to bridge this conversation.

Cross-Curricular Connections

This lesson connects to the following Florida middle school subject areas:

ELA (Reading Informational Text): the Student Reading is informational text written at grade level. Students practice identifying main ideas, supporting details, and the author’s purpose throughout. Vocabulary acquisition is explicit.

ELA (Communication and K-12 Expectations): the discussion in Activity 1 and the Reflection Questions exercise B.E.S.T. ELA’s K-12 Expectations for collaborative discussion (ELA.K12.EE.4.1 — collaborative techniques and active listening) and the grade-7 oral-presentation benchmark (ELA.7.C.2.1).

Mathematics: the AI’s inability to reliably do math is a concrete connection point. Discuss with math teacher: what makes AI different from a calculator? What kinds of math does AI struggle with? Why?

Science: the AI demonstration prompt about photosynthesis (in the No-Install version) connects directly to middle school life science. Coordinate with science teacher for terminology consistency.

Social Studies / Civics: the verification habit is the same skill students use in evaluating historical and contemporary sources. Coordinate with social studies teacher to reinforce that “check against real sources” is a habit they apply across subjects.

Parent Communication Notes

Most families will not need parent communication beyond the standard Parent / Guardian Letter sent before the curriculum begins. Specific situations that warrant additional communication:

A student becomes uncomfortable with the AI demonstration. Some students have anxiety about AI from media coverage. After class, send a brief note to the parent: “[Student] participated in our first AI literacy lesson today. They seemed [uncertain/uncomfortable] during the AI demonstration. The curriculum’s

approach is calm and grounded — we focus on understanding AI as a tool, not on AI as something to fear. Happy to discuss if you'd like."

A student is excited and wants to do more at home. Send the Foundation Edition reference (*The Open Source Student* v1.4 Chapter 1) and a brief note: "[Student] showed strong engagement today and is interested in going deeper. The Foundation Edition of the curriculum (which is suitable for adult or high-school readers) covers similar material in more depth. They may enjoy reading Chapter 1."

A parent asks specifically what AI tools were used in class. Be specific. "We demonstrated [tool name] on my classroom screen. Students did not interact with the AI directly. The Data Privacy Statement in the curriculum's Implementation Readiness Packet documents how AI is used in the curriculum."

Closing the lesson well

The transition from Lesson 1 to Lesson 2 is critical. End Lesson 1 with intentional setup for what comes next.

The setup language to use: "Today we learned that AI hallucinates — it makes up information that sounds true but isn't. Tomorrow we're going to take that seriously. We're going to take an actual AI answer and check it against real sources. By the end of tomorrow's lesson, you'll have a habit that protects you from the most common way students get into trouble with AI."

This frames Lesson 2 as the practical follow-through to Lesson 1's most important content. Students who are now thinking "okay, AI hallucinates — but what do I actually do about that?" are ready for Lesson 2.

Lesson 2: Verifying AI Answers With Sources

Teacher Edition

Lesson at a Glance

Lesson length: 45 minutes core instruction + 25 minutes student independent work (Activity 2 is the major artifact-producing work)

Learning objectives. By the end of this lesson, students will be able to: 1. Explain why verification matters when using AI 2. Apply a five-step verification workflow to a real AI claim 3. Distinguish between primary, secondary, and other source types and judge appropriateness for verification purposes 4. Produce a documented Source Verification Log entry that demonstrates verification practice 5. Connect the AI verification habit to the broader civic skill of source evaluation

Florida standards addressed: - **FL B.E.S.T. CS (grades 6–8):** Emerging Technologies (SC.7.ET.2.1), **Programming and Software Engineering — model limits (SC.8.PE.3.3, primary substance fit)**, Personal Health and Safety (SC.7.HS.1.4, SC.7.HS.1.6), Communication and Collaboration (SC.7.CC.2.2) - **Florida Civics:** SS.7.CG.2.9 (primary, full development — analyzing media for bias, symbolism, propaganda); SS.7.CG.2.8 (introduced — media/individuals/interest groups influencing government) - **B.E.S.T. ELA (grades 6–8):** Reading Informational Text — Argument (ELA.7.R.2.4), Communication — Research (ELA.7.C.4.1, primary fit), Vocabulary (ELA.7.V.1.1), K-12 ELA Expectations — Collaborative Discussion (ELA.K12.EE.4.1)

(See the Standards Alignment Matrix for benchmark-level alignment.)

Materials needed: - Teacher device with browser access (for AI demonstration in Activity 1, OR for displaying prepared example AI outputs) - Class display - Student copies of Lesson 2 Student Edition (printed or digital) - Source Verification Log Template (in Student Edition; can also be reproduced as separate worksheet) - Internet access for student devices during Activity 2 (for source-finding); paper alternative possible if a single shared computer is the only option - A pre-selected AI claim per student or per small group for Activity 2 (educator prepares 5–10 AI claims in advance to choose from)

Google Classroom quick access: Google Classroom Upload Pack — setup guide, copy/paste assignment posts, upload manifest, and Bellringer / Warm-Up pacing options.

Prerequisites: Lesson 1 (What AI Is and Is Not). The hallucination concept from Lesson 1 is essential context for this lesson.

Vocabulary introduced: Verification, Source, Credibility, Primary source, Secondary source, Citation

Portfolio artifact produced: Source Verification Log entry (Lesson 2 contribution to curriculum portfolio)

Assessment: Source Verification Log entry, evaluated with rubric below

---

Background for the Teacher

Lesson 2 is where the curriculum's verification habit becomes a practiced skill. Lesson 1 established that AI hallucinates; Lesson 2 gives students the workflow to do something about it.

The lesson's core move is structural: students learn that verification is not "I asked a teacher and they said it was fine" — it is a workflow with steps, sources, and documentation. The Source Verification Log artifact exists in part because the structure of "log it" is what makes the habit stick. Students who casually verify in their heads tend to drift away from the habit; students who produce documented log entries tend to keep doing it.

A few important framings:

Verification is a civic skill, not just an AI skill. This is the lesson where the curriculum's two pillars (AI literacy + civic technology literacy) most visibly converge. The same skill that protects students from AI hallucination protects them from misinformation, propaganda, and manipulation in any source type. The Florida civic standards that center on analyzing source bias and media influence (SS.7.CG.2.9 most directly, with SS.7.CG.2.8 supporting) are the same standards this lesson exercises. Note: SS.7.CG.2.10, often loosely described as the "source evaluation" benchmark, is in fact about the *process for citizens to address state or local problems* — that benchmark anchors Lesson 5, not Lesson 2.

Sources have a hierarchy. Students often need explicit teaching about why some sources are more reliable than others. The hierarchy in the reading (primary > reputable secondary > official > general reference > less reliable > almost never) is intentionally simplified for grade-level accessibility. Use the hierarchy as a discussion starter rather than as dogma; real source evaluation involves judgment, not just type-matching.

Two independent sources is the threshold. A single source can be wrong; multiple sources that all trace back to the same original could all be wrong together. Two independent sources — each with its own basis for the claim — substantially increases confidence. This is a real research methodology principle that students will use throughout their education.

Modern AI improvements are real but do not eliminate the need to verify. The lesson explicitly acknowledges that some AI systems use grounding, retrieval, and tool integration to reduce hallucination. This honesty matters — students who use modern AI tools have probably noticed that AI is better than it was. The lesson's response is "yes, AI is improving, and verification still matters." This is the calibrated position.

The lesson's structural goal is that by the end, every student has produced one Source Verification Log entry that demonstrates they can do verification. The habit only forms through practice; this lesson is the first practice.

Lesson Procedure

Total time: 45 minutes core instruction. Independent work afterward: 25 minutes (most of it spent on the Source Verification Log artifact).

Timing note for first delivery: Times are targets, not requirements. First-time delivery often runs 5–10 minutes long because students need help distinguishing a claim from evidence and choosing verification sources. If pacing gets tight, model one source together and assign the second source check as independent work.

Bellringer / Warm-Up Options (choose one)

Use one option at the start of class. If pacing is tight, use Option A and treat it as the first three minutes of the Opening rather than an added activity. Display this exact warm-up claim for all options: “Dolphins are fish because they live in water.”

Option A — Fast Start (3 minutes): Post: “Dolphins are fish because they live in water.” Students identify the exact claim that needs checking. Expected claim: dolphins are fish because they live in water. Use one response to introduce the word “verification.”

Option B — Standard Warm-Up (5 minutes): Post: “Dolphins are fish because they live in water. 1. What exact claim would you need to check? 2. What kind of source would you use to verify it? Examples: science textbook, reputable aquarium or marine biology website, encyclopedia, or government wildlife page. 3. What would count as enough evidence to trust your answer?” Have students compare sources they would use, then transition: “Today we are going to make that checking process systematic.”

Option C — Extended Launch (6-7 minutes): Post: “Dolphins are fish because they live in water. 1. What exact claim would you need to check? 2. What kind of source would you use to verify it? 3. What would count as enough evidence to trust your answer? 4. If two websites disagreed about this claim, what would you do next?” Use this to introduce independent sources and why one search result is not enough.

Opening (5 minutes)

Recap from Lesson 1: “Yesterday we learned that AI hallucinates. Who can explain in their own words what that means?” Get one or two student responses. Reinforce: AI sometimes produces wrong information in the same confident voice as right information; AI usually cannot tell the difference; this is built into how LLMs work and is a persistent risk that modern improvements reduce but do not eliminate.

Bridge: “Today we are doing something about it. We are learning a workflow called verification — checking AI claims against real sources. By the end of class, every one of you will have produced one Source Verification Log entry. This is a habit you will use for the rest of your life.”

State the lesson goal: “By the end of class, you will be able to take an AI answer that contains a factual claim, find real sources to check whether the claim is accurate, and document what you learned. You will also see how this same skill applies to anything else that comes at you as fact — news, social media, things people tell you.”

Direct Instruction (12 minutes)

Walk students through the Student Reading. Read aloud or have students read silently while you walk; choose based on your class’s reading patterns.

Key sections to emphasize:

“Why You Cannot Just Trust AI” — the Lincoln-quote example is concrete and memorable. Pause after this paragraph and ask: “Has anyone here had something like this happen — used an AI claim and later found out it was wrong?” If anyone has, briefly hear it. If not, that is fine; the example does the work.

“The Verification Workflow” — read all five steps. After each step, briefly explain the rationale. The five steps will be the structure of Activity 2; getting them clear here saves time later.

“What Sources Are Worth Checking Against” — the source hierarchy is the most important reference content in the lesson. Pause and discuss: “Why is a primary source more reliable than a secondary source?” “Why might a Wikipedia article be useful but not the final stop?” Make sure students understand that this is judgment, not a strict ranking.

“Why This Matters for Civic Literacy” — this paragraph makes the curriculum’s two-pillar framing visible. Read it aloud. Pause: “What other kinds of information besides AI do you encounter where this same skill would help?” Get 2–3 student responses. The point is that verification transfers beyond AI.

“What ‘Reduces, Not Eliminates’ Means” — for tech-savvy students who will say “but ChatGPT has search now,” this paragraph honors that observation while preserving the verification habit. Reinforce: AI has gotten better; verification still matters.

Activity 1: Verification Demonstration (10 minutes)

Demonstrate the verification workflow on a real AI response. The class watches. There are two paths depending on whether you have live AI access:

Path A — Live AI access available: - Type a query that will produce a response with verifiable claims. Suggested queries: - “Tell me three interesting facts about [a famous historical figure]” - “Explain a recent discovery in [a scientific field]” - “Give me a quote by [a real public figure] about [a topic]” - Project the response to the class - Walk through the five steps live: identify the specific claim, plan verification, find sources, compare, document

Path B — Prepared example outputs (no live AI): - Use the prepared example AI response from the No-Install Adaptation Notes section below - Display on class screen - Walk through the five steps using the prepared materials

In either path, do at least one verification where the AI is correct AND one where the AI is incorrect (the Einstein quote example from Lesson 1’s No-Install section is a reliable hallucination demonstration). Showing both outcomes is essential — students should not leave with the impression that verification always finds AI wrong (which would discourage AI use) or always finds AI right (which would devalue verification).

After the demonstration, ask: “What did you notice? What surprised you? How would you do this for your own AI claim?”

Activity 2: Student Verification (independent work, 18+ minutes)

This is the artifact-producing core of the lesson. Each student (or pair, depending on your setting) produces one Source Verification Log entry on an AI claim.

Setup (3 minutes): - Distribute the Source Verification Log Template (in the Student Edition) - Distribute or assign AI claims for verification — see "AI claims for verification" in the No-Install Adaptation Notes for prepared examples; alternatively, generate fresh ones during Activity 1 from your live AI access - Ensure each student has internet access for source-finding (or arrange paper alternative — see Differentiation Notes)

Student work (15 minutes): Students work independently or in pairs. Circulate. Specific things to look for:

- Are students writing the AI claim down exactly, or paraphrasing? (They should write it exactly; paraphrasing during transcription hides what was actually claimed.)
- Are students checking against truly independent sources, or two pages that obviously copy from each other? (Push toward independent sources.)
- Are students documenting honestly when the AI was right? (Some students try to make the AI look bad; gently correct: "Document what you found, not what you wished you found.")
- Are students stopping at one source? (Push for two; explain why.)

Wrap-up (closing 2 minutes of activity time): - Two or three students share what they found - Acknowledge: this is a real habit, you have just done it once, you will do it many more times across this curriculum and beyond

Closing (5 minutes)

Bring class back together. Highlight one specific student finding (with permission) — particularly powerful if a student found a hallucination during their verification.

Preview Lesson 3: "Tomorrow we take this verification skill and expand it. Instead of checking one AI claim against sources, we will look at how the same topic gets covered across many different kinds of sources — AI, news, advocacy organizations, social media. We will look for bias, missing context, and the patterns of how information actually moves through the world."

Direct students to Reflection Questions and Self-Check (in Student Edition) to complete after class or as homework.

Student Reading — Reproduced with Teacher Annotations

> **In Lesson 1, you learned that AI hallucinates** [TEACHER NOTE: Direct callback to Lesson 1's most important concept. If a student needs the recap, this is the moment.] **— that it sometimes produces wrong information in the same confident voice it uses for right information, and that it usually cannot tell the difference. Today's lesson is about what to do about that.**

> **The answer is short: verify.** [TEACHER NOTE: The lesson's central claim. Underline. Emphasize.]

Here is a concrete example. Imagine you ask an AI for a quote by Abraham Lincoln. [TEACHER NOTE: This example is concrete and memorable. Most students will quickly grasp how the failure mode works. The Lincoln example also generalizes — AI invents quotes from many real historical figures.]

Verification is a workflow, not a single check. [TEACHER NOTE: The "workflow" framing matters. Students who think of verification as "I'll just glance at one search result" do not really verify. The five-step structure is what makes verification reliable.]

Step 1: Identify the specific claim. [TEACHER NOTE: Underrated step. Students often try to "verify a paragraph" rather than identifying the specific claims within it. Demonstrate explicitly: "There are 3 claims in this AI response. Which one are we checking?"]

Step 2: Decide what counts as verification. [TEACHER NOTE: Different claims need different evidence. A historical date needs a historical source. A scientific claim needs a scientific source. A quote needs a confirmed primary source. Use Activity 1 to model this judgment.]

Step 3: Find at least two independent sources. [TEACHER NOTE: The "two independent" rule is research methodology. One source could be wrong. Two sources that obviously copy from each other are effectively one source. Two independent sources, each with their own basis for the claim, gives much higher confidence.]

Step 4: Compare the AI claim to the sources. [TEACHER NOTE: Students often want a binary "AI was right / wrong." The reality is messier — fully correct, mostly correct (small errors), partly correct (some pieces match), mostly wrong, fully invented. The Source Verification Log template provides this gradient.]

Step 5: Document what you found. [TEACHER NOTE: Documentation is what makes verification stick. Without it, students "verify" in their heads and forget. The Source Verification Log artifact exists for this reason. Treat the documentation as part of the workflow, not as a separate task.]

Not all sources are equally reliable. [TEACHER NOTE: The source hierarchy paragraph. Spend time here. The hierarchy is judgment-based, not strict. Wikipedia is genuinely useful — but as a starting point for finding citations, not as the citation itself.]

The verification habit is not just about AI. [TEACHER NOTE: This paragraph is the bridge to civic literacy. The same skill protects students from any kind of misinformation — news, social media, advocacy claims, things friends repeat. This is why source-bias evaluation appears in Florida civics standards (SS.7.CG.2.9 — analyzing media for bias/symbolism/propaganda, primary; SS.7.CG.2.8 — media/individuals/interest groups influencing government, supporting) and in B.E.S.T. ELA research standards (ELA.7.C.4.1 — research drawing on multiple reliable and valid sources). Note: SS.7.CG.2.10 is sometimes loosely cited here, but its actual text is about citizens addressing state or local problems — that benchmark anchors Lesson 5.]

Modern AI systems try to reduce hallucination. [TEACHER NOTE: Honest acknowledgment that AI is improving. Tech-savvy students will appreciate this honesty. The point: AI improvement is real, the verification habit is still necessary.]

Vocabulary Teaching Strategies

Six terms today (verification, source, credibility, primary source, secondary source, citation). Most students have heard these in social studies; the lesson sharpens their meaning in an information-evaluation context.

Pre-teach during opening: ask the class “what does it mean to verify something?” Get a definition or two. Refine if needed: “Verifying means checking whether something is actually true by comparing it against a source you can trust.”

Primary vs. secondary distinction: this is the trickiest pair for middle school students. Use a concrete example: “If we want to know what Abraham Lincoln actually said, the primary source is a written record of him saying it — a letter, a speech transcript, a contemporary newspaper account. A textbook talking about Lincoln’s quotes is a secondary source — useful, but if the textbook quotes Lincoln, you might want to find the primary source it cites to be sure.” Most students grasp this with one or two examples.

For ELL students: “verificación” is a direct cognate. “Source” maps to “fuente.” “Credibility” maps to “credibilidad.” Most of today’s vocabulary has Spanish cognates, which is helpful.

Vocabulary in use during Activity 2: the Source Verification Log template uses all six terms. By the end of the lesson, every student should have used the vocabulary in their own writing.

Differentiation Notes

Below grade level readers

- Read the verification workflow aloud while students follow along
- Provide a one-sentence version of each step on the board
- For Activity 2, pair below-level readers with on-level partners
- Allow oral verification (student finds the source and explains aloud what it says, rather than writing the comparison)

Above grade level readers

- Extension reading: Foundation Edition Chapter 2 (the longer treatment of why verification matters)
- Have advanced students do their verification on a more complex AI claim (a multi-paragraph response with multiple verifiable claims)
- Have advanced students serve as verification helpers for classmates struggling with the workflow

English Language Learners (ELL)

- Most vocabulary has Spanish cognates; flag them for ELL students
- Allow Spanish-language sources where appropriate (Spanish Wikipedia, Spanish-language news, etc.)
- Provide sentence frames for the Source Verification Log entry: “The AI claimed _____. Source 1 says _____. Source 2 says _____. So the AI was _____ because _____.”
- For newcomer ELL students, focus on a simpler verification (e.g., a single date claim) rather than a complex one

IEP / 504 — ADHD / Executive Function

- Use the five-step workflow as an explicit checklist; check off each step as completed
- Break the 18-minute Activity 2 into shorter visible segments (5 min identify claim, 5 min find sources, 5 min compare, 3 min document)
- Provide pre-formatted Source Verification Log template with sections clearly demarcated

IEP / 504 — Reading or Language Processing

- Audio of the Student Reading (TTS or recorded)
- Allow oral or recorded Source Verification Log entry instead of written
- Larger-font printable template

IEP / 504 — Autism Spectrum

- Provide the lesson structure visually before starting (the same five-component lesson structure students saw in Lesson 1)
- For Activity 2, the systematic five-step workflow may be a strength — autistic students often handle structured procedures well; explicitly acknowledge this if helpful
- If group work in Activity 1 demonstration is overwhelming, allow private observation with later check-in

Anxiety or Emotional Regulation

- Reassure: there is no “right” outcome to verification. Both “the AI was correct” and “the AI was wrong” are legitimate findings.
- The Source Verification Log is graded on process (did the student do the workflow), not outcome (what they found)
- Reflection Questions can be answered privately

Limited Technology Access

- For Activity 2, students without internet can verify against printed sources (textbook, encyclopedia, library books)
- Provide a small set of pre-printed source materials educators can hand out for verification (a textbook excerpt, an encyclopedia entry, a printed government document)
- The Source Verification Log can be completed on paper

Gifted students

- Have them verify a claim where the answer is genuinely unclear or contested — “what is the average global temperature change since 1900?” requires nuanced source evaluation
- Have them research the verification methodologies that journalists use (fact-checking organizations like PolitiFact, FactCheck.org)
- Connect to Foundation Edition Appendix D (AI Agent Platforms) — modern retrieval-augmented systems try to automate verification; what would it take to make that reliable?

Assessment Rubric: Source Verification Log Entry

The Source Verification Log entry is the portfolio artifact for Lesson 2. Evaluate primarily on whether the verification workflow was followed, not on whether the AI was right or wrong.

Criterion	Developing	Proficient	Exemplary
Specific claim identified	Vague claim or paraphrase	Specific claim quoted exactly from AI	Specific claim plus brief explanation of why this claim matters
Verification plan stated	No plan; jumped to searching	Plan stated: what kind of source needed, where to look	Plan stated and reasoning given (e.g., “this is a historical claim, so I want a primary source from the Lincoln Presidential Library”)
Two sources found and documented	One source only, or sources not documented	Two sources with type and reference noted	Two independent sources of clearly different types (e.g., one primary, one reputable secondary)
Comparison shows real engagement	Vague “the AI was right/wrong”	Specific comparison: which parts match, which don’t	Nuanced comparison: identifies “mostly correct but with a small error” or “partly true with missing context”
Documentation completeness	Missing key fields	All template fields completed	All fields completed plus reflection on the process

Criterion	Developing	Proficient	Exemplary
Honesty about findings	Findings may be exaggerated either toward “AI right” or “AI wrong”	Honest report of what sources said	Honest report includes uncertainty: “I am not sure whether this counts as the AI being wrong or just incomplete”
What was learned	Generic (“I learned to verify”)	Specific takeaway from this verification (“I learned that AI quotes from real people are often invented”)	Specific takeaway connected to broader habit (“I am going to verify quotes specifically from now on, because this is the second time I’ve seen AI invent one”)

Sample Proficient artifact (for educator reference):

```
Date: May 14, 2026
Student: M.S.

AI Response Source: Classroom demonstration, ChatGPT

The Claim Being Verified: "Abraham Lincoln said, 'In the end, it's not the years
in your life that count. It's the life in your years.'"

Why This Claim Matters: I want to use this quote in my essay about Lincoln, so I
need to know if he actually said it.

Verification Plan: I want to find a primary source – a real Lincoln document or
the Lincoln Presidential Library website – or a reputable Lincoln biography that
quotes the actual source.

Source 1:
- Type: official source
- Reference: Lincoln Presidential Library website (https://...)
- What this source says: I searched for the quote on the Lincoln Presidential
  Library site. The quote does not appear in any of their collections.

Source 2:
- Type: reputable secondary source / fact-check
- Reference: QuoteInvestigator.com article on this exact quote
- What this source says: This quote is widely misattributed to Lincoln. The actual
  origin is uncertain but appears to be from a 1947 magazine, decades after
  Lincoln died. Lincoln did not say it.

Comparison: Fully invented – AI claim is not supported by any reliable source

What I Found: The AI confidently attributed this quote to Lincoln. Lincoln did
not say it. The quote comes from a 20th-century source and has been misattributed
to many famous people, not just Lincoln.

What I Learned About AI From This Verification: AI seems to be especially likely
to invent quotes from famous historical figures. I'm going to verify any quote I
plan to use, regardless of who it's attributed to.
```


This sample would receive Proficient or Exemplary marks depending on the depth of the “what I learned” reflection.

Answer Key / Scoring Notes

Use these notes to calibrate student verification work. The lesson is graded on process quality, not on whether the AI claim turned out true or false.

Bellringer / Warm-Up answer key: - Exact claim: **“Dolphins are fish because they live in water.”** - Correct verification outcome: dolphins are **mammals**, not fish. Living in water is not enough to classify an animal as a fish. - Reliable verification sources include a science textbook, aquarium or marine biology organization, encyclopedia, NOAA or other government wildlife page, or teacher-approved science source. - Evidence that is enough: a source explaining that dolphins breathe air, nurse young, are warm-blooded mammals, and belong to marine mammal groups.

Activity 1 demonstration expected outcomes: - Students should identify one specific factual claim from the AI response before searching. - They should say what kind of source would verify that claim. - They should compare the AI wording to the source wording, not merely announce “right” or “wrong.” - If the AI is partly correct, students should record the nuance. “Mostly correct but missing context” is a valid outcome.

Source Verification Log minimum proficient standard: - Exact claim copied or quoted. - A verification plan that matches the claim type. - Two independent sources documented with source type and reference. - A comparison that identifies what matches and what does not. - A final statement of what the student learned about the claim and about AI reliability.

Self-Check answer key: 1. Verification matters because AI can produce confident but false or unsupported claims; verification protects student work and civic judgment. 2. A primary source comes directly from the original event, person, organization, record, or data source. A secondary source explains, interprets, summarizes, or reports on primary sources. 3. Two independent sources reduce the chance that one source is mistaken, copied, biased, or incomplete. Two pages repeating the same source do not count as truly independent. 4. Usually reliable source types include official government pages, primary documents, textbooks, academic or expert sources, reputable news reporting, and established encyclopedias used as starting points. Usually weaker sources include unsourced social media posts, anonymous blogs, copied quote pages, AI output by itself, and advocacy claims with no evidence. 5. Verification is not only for AI. It also applies to news stories, social media, advocacy claims, rumors, friend-to-friend claims, and civic information.

Common grading mistakes to avoid: - Do not give a high score for finding a correct answer if the student did not document the workflow. - Do not penalize a student for finding that AI was correct; correct AI output is common and worth documenting. - Do not accept “I Googled it” as documentation. The student must name the source used and what it said.

Extension Options

For students who finish early or want to go deeper:

Extension 1: Verify a multi-claim response. Take an AI response with several factual claims and verify each one separately. Document which were correct, which had errors, which were invented.

Extension 2: Compare verification across AI services. Take the same factual question, ask two different AI services, and verify both responses. Were they the same? Different? What does this tell us about AI reliability?

Extension 3: Investigate hallucination patterns. Look up published research or journalism on AI hallucination types. What kinds of hallucinations are most common? What categories of claim is AI most likely to get wrong?

Extension 4: Research professional verification. What do professional journalists do for verification? Look up the methodology of one or two fact-checking organizations (PolitiFact, FactCheck.org, Snopes). How does their process compare to what we did today?

No-Install Adaptation Notes

If you do not have live AI access for Activity 1, use the prepared materials below.

Prepared AI claims for verification (Activity 2 source material)

These are pre-prepared AI claims for student verification practice. Mix verifiable-correct, verifiable-incorrect, and verifiable-invented claims so students experience the full range of outcomes.

Claim Set A (likely true; verification confirms): 1. “The Roman Republic ended in 27 BCE when Augustus became the first Emperor.” 2. “Photosynthesis converts carbon dioxide and water into glucose using energy from sunlight.” 3. “Florida became a U.S. state in 1845, the 27th state to join the Union.”

Claim Set B (likely incorrect with subtle error; verification reveals the error): 1. “The Roman Republic ended in 31 BCE.” [Correct year: 27 BCE; 31 BCE was the Battle of Actium] 2. “Photosynthesis happens in the cell’s mitochondria.” [Wrong organelle; happens in chloroplasts] 3. “Florida became a U.S. state in 1845, the 28th state to join the Union.” [27th, not 28th]

Claim Set C (likely invented; verification reveals fabrication): 1. “Mark Twain wrote, ‘The two most important days in your life are the day you are born and the day you find out why.’” [Widely misattributed; not Twain] 2. “Eleanor Roosevelt said, ‘No one can make you feel inferior without your consent.’” [Widely attributed to her but not actually a documented quote of hers] 3. “Albert Einstein said, ‘Education is what remains after one has forgotten what one has learned in school.’” [Misattribution; actually from a different source]

Distribute one claim per student, varying across the three sets so the class collectively experiences different outcomes.

Verification demonstration without live AI

Use Claim B-1 or C-1 as your demonstration claim. Walk through the five-step workflow:

1. Identify the claim (read it aloud, write it on the board)

2. Plan verification (where would we look? — historical source for B-1; primary source documenting Twain's actual writings for C-1)
3. Find sources — model the search; for the prepared demonstration, you can have a printed reference page ready
4. Compare — show the AI claim and the source side by side
5. Document — fill in a Source Verification Log template on the board as you go

The demonstration takes about 8 minutes and reaches the same learning objective as live AI demonstration.

Common Misconceptions

"If I can't find it on Google, it must be made up." Not necessarily. Some real sources are not indexed by general search engines, are paywalled, or require library access. For some claims, "I can't easily verify" means the verification process needs to escalate (ask a librarian, use academic databases) rather than concluding the AI is wrong.

"If the first source I found agrees with the AI, the AI must be right." Not necessarily. The first source may itself be wrong, may be parroting the AI's same misinformation, or may be incomplete. The two-source rule exists for this reason.

"Verification is too much work for normal use." It can feel that way at first. With practice, the workflow gets faster — most verifications take 2–5 minutes, not 30. The lesson's framing: do verification when the fact actually matters; for casual exploration, you may not need formal verification.

"AI got it right this time, so it's fine." AI being right one time tells you nothing about whether it will be right next time. The verification habit is about the kind of work you are doing (work where being wrong matters), not about the AI's track record. Even AI that is usually right occasionally hallucinates.

"Wikipedia is always wrong / Wikipedia is always right." Both are wrong. Wikipedia is a useful starting point — its citations point to real sources you can then verify. Wikipedia articles vary in quality; well-developed articles on widely-edited topics are usually accurate, while obscure topics or recent events may be less reliable. Use Wikipedia to find sources, then verify those sources directly.

Cross-Curricular Connections

Social Studies (Civics): SS.7.CG.2.9 (analyzing media and political communications for bias, symbolism, and propaganda) is the most direct grade-7 alignment, with SS.7.CG.2.8 (media/individuals/interest groups influencing government) as supporting context. Note: SS.7.CG.2.10 — sometimes loosely described as "source evaluation" — is in fact the *citizens addressing state or local problems* benchmark and anchors Lesson 5 rather than this lesson. Coordinate with the social studies teacher: today's lesson is the same source-bias evaluation work students will do on historical sources in U.S. History or on political communications in civics.

ELA (Research): B.E.S.T. ELA Communication standards (ELA.7.C.4.1) emphasize research and inquiry, including locating multiple sources, evaluating credibility, and documenting sources. Today's lesson is the foundational research methodology those Communication-strand standards assume. (Note: B.E.S.T. ELA uses "C" for Communicating rather than "W" for Writing, which is the strand prefix in some other state frameworks.)

Science: scientific evidence relies on primary sources (peer-reviewed research) verified across multiple studies. The two-source rule and primary-source preference apply directly to scientific reasoning. Coordinate with the science teacher to reinforce.

Mathematics: for AI claims about statistics or numerical relationships, students may need to evaluate the math underlying the claim, not just whether sources agree. Coordinate with the math teacher about how to evaluate quantitative claims.

Parent Communication Notes

Most parents will appreciate this lesson without specific communication. The verification habit is a clearly responsible practice that aligns with what most parents already want from school — students who think critically rather than accept claims at face value.

Specific situations that warrant additional communication:

A student finds something genuinely surprising during verification (e.g., that a quote they had heard for years is misattributed). This is a teachable moment worth sharing: "[Student] verified an AI claim today and discovered that a quote often attributed to [person] is misattributed. They asked great questions about how widespread misinformation can be. You may enjoy continuing the conversation at home."

A student becomes worried about how much misinformation is "out there." Normalize: this is how many adults feel when they first develop the verification habit. Reassure: the verification habit is the response. They are now better equipped to handle the information environment than most adults are.

A parent asks for the verification framework to use at home. Share the five-step workflow from the lesson. It applies to any source, not just AI.

Closing the lesson well

Lesson 2 establishes the verification habit. Lesson 3 expands it to a wider information landscape. The transition matters.

End-of-lesson language: "Today you practiced verification on one AI claim. Tomorrow we are going to take this same skill and apply it to something bigger — looking at how information moves through different sources, where bias appears, what gets left out. The verification habit you started today becomes the foundation for everything we do for the rest of this curriculum, and for the kind of citizen you grow up to be."

This frames Lesson 3 as the natural extension and the curriculum's broader civic purpose as the larger arc.

Lesson 3: Media Bias, Claims, and Digital Persuasion

Teacher Edition

Lesson at a Glance

Lesson length: 45 minutes core instruction + 25 minutes student independent work

Learning objectives. By the end of this lesson, students will be able to: 1. Apply the six-question framework (claim, evidence, source, audience, missing context, bias) to any information source 2. Compare how the same topic is treated across multiple source types (AI, news, advocacy, social media) 3. Identify bias in sources, including sources they instinctively trust 4. Recognize missing context and explain how it shapes meaning 5. Produce a documented Media Analysis demonstrating cross-source critical reading 6. Articulate the connection between source-evaluation skills and civic literacy

Florida standards addressed: - **FL B.E.S.T. CS (grades 6–8):** Technological Impact — SC.7.TI.1.3 (“Identify how media is used to influence information,” strong primary fit); Personal Health and Safety — SC.7.HS.1.7, SC.8.HS.3.2 (“Analyze how digital media and communication influence behavior”); Communication and Collaboration - **Florida Civics:** SS.7.CG.2.9 (primary, full development — analyzing media for bias, symbolism, propaganda); SS.7.CG.2.8 (continued — media/individuals/interest groups influencing government); SS.8.CG.2.4 (grade 8 civic-participation history extension, taught within M/J U.S. History #2100010) - **B.E.S.T. ELA:** Reading Informational Text — Argument (ELA.7.R.2.4), Comparative (ELA.7.R.3.3), Rhetoric (ELA.7.R.3.4); Communication — Argument (ELA.7.C.1.3 — note B.E.S.T. uses C.1 prefix, not W.1); Vocabulary (ELA.7.V.1.1); K-12 ELA Expectations (ELA.K12.EE.4.1 collaborative discussion, ELA.K12.EE.6.1 voice and tone)

(See the Standards Alignment Matrix for benchmark-level alignment.)

Materials needed: - Teacher device with browser access (for AI access during demonstration; live AI optional with prepared example AI outputs as substitute) - Class display - Pre-selected topic for class demonstration of the Information Integrity Lab Mini - Pre-prepared sample sources for demonstration (see No-Install Adaptation Notes) - Student copies of Lesson 3 Student Edition - Internet access for student devices during Activity 2 (or printed source packets if devices unavailable) - Media Analysis Template (in Student Edition)

Google Classroom quick access: Google Classroom Upload Pack — setup guide, copy/paste assignment posts, upload manifest, and Bellringer / Warm-Up pacing options.

Prerequisites: Lessons 1 and 2 complete. The verification habit from Lesson 2 underlies today’s analysis work.

Vocabulary introduced: Bias, Claim, Evidence, Audience, Missing context, Persuasion, Information Integrity

Portfolio artifact produced: Media Analysis (Lesson 3 contribution)

Assessment: Media Analysis evaluated with rubric below

Background for the Teacher

Lesson 3 is the curriculum's most ambitious lesson conceptually. It asks students to do something hard: read multiple sources critically at the same time, identify patterns and gaps across them, and recognize bias in sources they trust as well as sources they don't. This is the work of civic literacy at the middle-school level.

Three framings are important to get right:

Bias is universal. Students sometimes arrive at this lesson with a "bias means lying" association — and a sense that "my side" is unbiased while "the other side" is. The lesson reframes: every source has perspective; the question is what the perspective is and how it shapes what is said. This applies to sources the student likes, sources their parents like, sources the school assigns, and sources nobody likes. Frame this as intellectual honesty, not as cynicism.

Missing context is harder to see than direct misinformation. Direct misinformation is easier to call out — the source said something that is wrong. Missing context is sneakier — the source said something that is technically true but leaves out the surrounding information that would change the reader's interpretation. Most modern misinformation works through missing context, not outright lies. Teaching students to look for what is left out is one of the highest-leverage skills in this lesson.

The lesson is the bridge to civic literacy. Florida civics standards SS.7.CG.2.9 (analyzing media and political communications for bias, symbolism, and propaganda — primary), SS.7.CG.2.8 (impact of media, individuals, and interest groups on monitoring and influencing government — supporting), and SS.8.CG.2.4 (grade 8 civic-participation history extension, taught within M/J U.S. History #2100010) ask students to evaluate sources on public issues and the media-and-democracy patterns those sources sit inside. This lesson is the methodology that makes those standards practical. A student who can do the Information Integrity Lab Mini on a topic can do it on any civic question they encounter.

The lesson is intentionally non-partisan in its examples and topics. The skill is the goal; specific political conclusions are not. When students discover bias in sources that match their own views, that is a successful learning outcome — not a problem to manage.

Lesson Procedure

Total: 45 minutes core instruction + 25 minutes student independent work.

Timing note for first delivery: Times are targets, not requirements. First-time delivery often runs 5–10 minutes long because the six-question framework is new and students need practice naming missing context. If pacing gets tight, reduce the class demonstration to three sources and let students finish the synthesis paragraph independently.

Bellringer / Warm-Up Options (choose one)

Use one option at the start of class. If pacing is tight, use Option A and treat it as the first three minutes of the Opening rather than an added activity. Display this exact warm-up sentence for all options: “This new school policy will solve the problem for everyone.”

Option A — Fast Start (3 minutes): Post: “This new school policy will solve the problem for everyone.” Students identify the claim being made. Expected claim: the new school policy will solve the problem for everyone. Ask one student to read their claim statement aloud.

Option B — Standard Warm-Up (5 minutes): Post: “This new school policy will solve the problem for everyone. 1. What claim is being made? 2. Who might be the audience? Examples: students, parents, teachers, administrators, school board members, or voters. 3. What important context is missing? Examples: what the policy is, what problem it addresses, who benefits, who disagrees, or what evidence supports it.” Have students share what context is missing. Use that to introduce the six-question framework.

Option C — Extended Launch (6-7 minutes): Post: “This new school policy will solve the problem for everyone. 1. What claim is being made? 2. Who might be the audience? 3. What important context is missing? 4. How could two people with different opinions use this sentence in different ways?” Use this to introduce audience, framing, and persuasion.

Opening (5 minutes)

Recap from Lesson 2: “Yesterday you verified an AI claim against real sources. You learned the verification workflow. Today we expand it. Instead of one AI claim against sources, we are going to look at one topic across many sources — AI, news, advocacy, social media — and see what each one is doing.”

Bridge: “By the end of class, you will be able to look at any information you encounter and ask six questions about it. This is one of the most important skills any citizen has.”

State the lesson goal as written in the Student Edition.

Direct Instruction (15 minutes)

Walk students through the Student Reading. The reading is longer and conceptually denser than Lessons 1 and 2; pace accordingly.

Key sections to emphasize:

“Information Doesn’t Travel Through One Pipe” — set up the lesson’s premise. The same topic gets discussed in many places by many actors with many purposes. No single source is “the truth”; all are interpretations.

“Information Has Always Been Messy. AI Just Made It Faster.” — important framing. Students sometimes blame AI for misinformation. The honest answer: people have always shaped information; AI accelerated the speed and scale, which is a real problem, but the underlying skills (verification, source evaluation, critical reading) are timeless.

“What to Look For in Any Source” — the six questions framework. This is the lesson’s most important content. Spend time here. Walk through each question with a concrete example. Use a familiar source type (e.g., a recent news headline students have seen) to demonstrate the framework before applying it formally

in Activity 1.

“What This Has to Do With Civic Life” — the bridge to civic literacy. The Florida civics standards reason for this lesson. Make the connection explicit so students understand why this is in their curriculum.

“A Note on Bias” — read this paragraph carefully. The reframing — every source has bias; the question is what kind, where, how it shapes the message — is one students need to hear directly. Some students will resist; let them work through the resistance during Activity 2.

Activity 1: Class Demonstration (10 minutes)

Pick a topic and lead the class through the six-question framework on three or four sources covering it. Topic recommendations (choose one):

- **Sea level rise impacts in Florida** (concrete, local relevance, plenty of source diversity)
- **Use of phones in middle schools** (immediately relevant to students’ lives)
- **Endangered species protection in your county** (local civic relevance)

For the chosen topic, gather (in advance): - An AI response on the topic - A news article from a recognized news outlet - An advocacy organization’s framing of the topic - (Optional) a social media post discussing the topic

For each source, lead the class through the six questions on the board: - What is the claim? - What evidence is offered? - What is the source? - Who is the audience? - What context is missing? - Where is the bias?

The demonstration takes about 8 minutes if you keep it disciplined. After analyzing 3–4 sources, do the synthesis question: “What does the pattern across these sources tell us? What does any single source leave out? What would a careful citizen need that no single source provides?”

Activity 2: Student Media Analysis (15 minutes core, plus independent work for completion)

Direct students to Activity 2 in the Student Edition. They will produce their own Media Analysis on a topic.

Topic selection: offer the suggested topics in the Student Edition or assign topics if appropriate. For Florida pilot adopters, topics with local relevance produce the most engaged work.

Source gathering (5 minutes): - AI response: assist students in obtaining one (live AI if available, prepared examples otherwise) - News article: students search a recognized news site (or you provide pre-selected articles from a Florida news source) - Advocacy or government source: students find one (or you provide) - Optional social media post: only if appropriate for your setting

Analysis (10+ minutes): - Students apply the six questions to each source - They write the synthesis paragraph - They complete the Media Analysis template

If students do not finish in class, they complete as homework.

During work, circulate and look for: - Are students naming bias in sources they instinctively trust? (Push them to do this — it is the hardest and most important part.) - Are students noticing missing context? (Often the most overlooked question; prompt: “What’s NOT in this source that should be?”) - Are students

producing real synthesis or just listing sources? (The pattern paragraph is what makes the artifact analytical rather than descriptive.)

Closing (5 minutes)

Bring class back. Ask 2–3 students to share one thing they noticed during their analysis — particularly bias in a source they expected to be neutral.

Acknowledge: "This work is genuinely hard. You will not get perfect at it in one lesson. The framework is now in your head; over time, you will use it without thinking."

Preview Lesson 4: "Tomorrow we shift from analyzing other people's information to thinking about how YOU use AI — privacy, ethics, academic integrity. The skills you have built so far (verification, source analysis) are about evaluating information from outside. Tomorrow is about how you use information well in your own work."

Student Reading — Reproduced with Teacher Annotations

> **In Lesson 2, you practiced verifying a single AI claim against real sources.** [TEACHER NOTE: Direct callback to Lesson 2. The verification skill underlies today's broader analysis.]

> **A typical topic in the modern world might be discussed: by an AI service... in news articles... by advocacy organizations... across thousands of social media posts...** [TEACHER NOTE: This list is the lesson's core insight. Students often think information comes from one place; the list makes the multi-source reality concrete.]

> **It is tempting to think that misinformation, bias, and persuasion are new problems caused by AI. They are not.** [TEACHER NOTE: This framing matters. Students who blame AI for misinformation will neither use AI well nor handle non-AI sources critically. The honest framing — older problems made faster by AI — produces better student outcomes.]

> **What to look for in any source.** [TEACHER NOTE: The six-question framework. The most important content in the lesson. Students should leave able to recite these six questions.]

> **1. What is the claim?** [TEACHER NOTE: Students often try to "evaluate" without first identifying what is being claimed. Force this step. Many sources hide their actual claim behind framing or rhetoric; naming

the claim is half the work.]

2. What evidence is offered? [TEACHER NOTE: Strong sources cite verifiable evidence. Weak sources offer vague references (“studies show”), emotional appeal, or unsupported assertion.]

3. What is the source? [TEACHER NOTE: Connect to Lesson 2 source-evaluation work. Who produced this? What’s their incentive? Is this a credentialed expert with no obvious stake, or an organization that benefits if the claim is true?]

4. Who is the audience? [TEACHER NOTE: Underrated question. Sources frame claims for their audiences. Understanding audience is key to understanding why a source emphasizes what it does.]

5. What context is missing? [TEACHER NOTE: One of the highest-leverage questions. Most modern misinformation works through missing context, not outright lies. Push students to ask: what’s NOT here?]

6. Where is the bias? [TEACHER NOTE: The “A Note on Bias” paragraph addresses this in depth. Students may need encouragement to name bias in sources they trust — this is the hardest version of the question.]

The reason source evaluation appears in Florida civics standards [TEACHER NOTE: The bridge to civic literacy. Make this connection explicit during direct instruction.]

A note on bias. [TEACHER NOTE: Read this paragraph carefully during direct instruction. The “every source has bias” framing is the lesson’s intellectual move. Students who resist this framing benefit from working through the resistance during Activity 2.]

Vocabulary Teaching Strategies

Seven terms today (more than previous lessons because the lesson is conceptually richer).

Pre-teach during opening: ask “what does bias mean?” Get definitions. Refine: “Bias is a systematic lean in how something is presented. It’s not always bad — everyone has perspectives — but unrecognized bias can shape what we believe without us knowing.”

Most important pair to distinguish: claim vs. evidence. Students often conflate these. A claim is what is asserted; evidence is what supports the claim. A source can have a strong claim with weak evidence, or a weak claim with strong evidence. The framework asks both questions for a reason.

Most subtle term: missing context. Many students have not thought of “what is left out” as a thing to evaluate. Use a concrete example: “If a source says ‘crime is up 20% this year,’ missing context might include: compared to what year? In what category? In what location? With what definition of crime?”

For ELL students: “bias” maps roughly to “sesgo” or “parcialidad.” “Claim” maps to “afirmación.” “Evidence” maps to “evidencia.” “Audience” maps to “audiencia.” Most cognates work; flag for ELL students.

Differentiation Notes

Below grade level readers

- Walk through the reading section by section, with check-ins
- Provide a simplified version of the six-question framework on a single visual card students can keep
- For Activity 2, simplify by reducing source count to 2 instead of 3

Above grade level readers

- Extension: have them analyze 4 or 5 sources instead of 3
- Extension: have them propose what would make each source’s coverage more complete
- Extension reading: Foundation Edition Chapter 14 covers ethics adjacent to bias; relevant for tomorrow’s lesson

English Language Learners

- Most vocabulary has Spanish cognates; flag them
- Pair with bilingual peers during analysis
- Allow Spanish-language sources for the Media Analysis (Spanish news, Spanish advocacy organizations); cross-language analysis is a particular strength for ELL students
- Provide sentence frames for the analysis: “The claim is _____. The evidence is _____. The source is _____.”

IEP / 504 — ADHD / Executive Function

- Use the six-question framework as an explicit checklist
- Break the analysis work into 4–5 minute segments per source
- Provide pre-formatted Media Analysis template with sections clearly demarcated

IEP / 504 — Reading or Language Processing

- Audio of Student Reading
- Allow oral analysis (student explains their reading of each source aloud rather than writing each section)
- Pair with on-grade peer for source reading

IEP / 504 — Autism Spectrum

- The systematic six-question framework may be a strength for many autistic students
- Allow private work during Activity 2 if group dynamics are overwhelming
- Provide explicit social expectations for the closing share-out

Anxiety or Emotional Regulation

- The bias content can feel charged; reassure that there is no “right” finding
- Make explicit: students will not be graded on what bias they find; they will be graded on whether they did the analysis honestly
- Allow students to choose a low-charge topic for their Media Analysis if a high-charge topic would be difficult

Limited Technology Access

- Provide pre-printed source packets (one packet per topic, 3–4 sources each) for students who cannot search online
- The Media Analysis can be completed entirely on paper

Gifted students

- Have them analyze a topic with genuinely contested expert disagreement (climate adaptation strategies, education policy options) rather than one with clear-cut sourcing
- Have them research how professional fact-checkers and media analysts approach this work
- Have them produce a comparison of how two different news outlets cover the same story

---

Assessment Rubric: Media Analysis

Criterion	Developing	Proficient	Exemplary
Topic specificity	Topic too broad to analyze	Topic specific enough for meaningful comparison	Topic specific to a real local or current issue
Number and variety of sources	Fewer than 3 sources, or all the same type	3 sources of at least 2 different types	3 or more sources spanning AI, news, advocacy/government, and (optionally) social media
Six-question framework applied to each source	Some questions skipped or answered superficially	All six questions answered for each source with substance	Six questions answered with specific evidence cited from each source
Bias identification	Bias named only in “obvious” sources	Bias identified in all sources, including ones the student may agree with	Bias identification is nuanced — distinguishes degree, type, and effect of bias

Criterion	Developing	Proficient	Exemplary
Missing context identification	“Missing context” answered vaguely or skipped	Specific missing context noted for each source	Missing context analysis includes “what would a citizen need that this leaves out?”
Synthesis paragraph	List of sources without pattern recognition	Pattern across sources articulated; agreement and disagreement named	Synthesis identifies what no single source provides; offers analytical insight beyond any individual source
Honest engagement with discomfort	Student avoids naming bias in trusted sources	Student names bias in trusted sources where it exists	Student reflects in writing on what was hardest about the analysis

Answer Key / Scoring Notes

Lesson 3 has no single answer key because students may analyze different topics and sources. The teacher-facing key is the **six-question framework** plus the quality indicators below.

Six-question framework key: 1. **Claim:** what the source wants the reader/viewer to believe. 2. **Evidence:** facts, examples, data, expert statements, documents, images, personal experience, or reasoning offered to support the claim. 3. **Source:** who produced the message and what role, expertise, incentive, or stake they have. 4. **Audience:** who the message is trying to reach or persuade. 5. **Missing context:** what a careful reader still needs to know before deciding what to believe. 6. **Bias:** the perspective, lean, selection pattern, emotional framing, incentive, or omission shaping how the information is presented.

Self-Check answer key: 1. Bias is a systematic lean or perspective in how information is selected, framed, emphasized, or omitted. It is not always dishonest, but it must be recognized. 2. The six questions are: claim, evidence, source, audience, missing context, and bias. 3. Every source has bias because every source is created by a person, organization, tool, or institution with a perspective, purpose, audience, limits, and choices about what to include or leave out. 4. Missing context means important information needed to understand the claim is absent. Example: “Crime is up 20%” is missing location, time period, crime category, data source, and comparison baseline. 5. These are civic-literacy skills because citizens need to evaluate public claims, media messages, advocacy arguments, official statements, and AI-generated information before forming opinions or acting in public life.

Common grading mistakes to avoid: - Do not require students to label one source “good” and another “bad.” The goal is to analyze how each source works. - Do not accept “this source is biased” as a complete answer. Students should say what kind of bias and how it shapes the message. - Do not treat government or trusted news sources as bias-free. They may be more accountable or evidence-based, but they still make choices about framing and emphasis.

Sample Proficient Artifact (for educator reference)

```
INFORMATION INTEGRITY LAB MINI — MEDIA ANALYSIS

Date: May 14, 2026
Student: J.R.

Topic: Cell phone restrictions in middle schools
```

```
Source 1: AI response (type: AI)
- Claim: Schools should limit phones during class because phones distract students,
  but students may still need phones for safety and communication.
- Evidence: The response mentions distraction, learning focus, and family contact,
  but does not cite any specific study or school policy.
- Source: AI system used in class; it summarizes common arguments but is not a
  primary source.
- Audience: A general student or teacher audience.
- Missing context: No local school policy, no data from our district, no student
  survey, and no research citation.
- Bias: Balanced-sounding, but it may smooth out conflict and make the issue seem
  simpler than it is.

Source 2: Local news article (type: news)
- Claim: More schools are restricting phone use because teachers report distraction
  and administrators want more focused classrooms.
- Evidence: Quotes from a principal and two teachers; mentions discipline referrals
  decreasing after one school changed its policy.
- Source: Local news reporter covering education.
- Audience: Parents and community members.
- Missing context: Few student voices; no full data table; little attention to
  students who use phones for medical or family reasons.
- Bias: Emphasizes adult and school-administration perspectives.

Source 3: Parent advocacy post (type: advocacy)
- Claim: Students should be allowed to keep phones because families need direct
  contact during emergencies.
- Evidence: Uses a parent story about not reaching a child quickly after an
  incident, but does not compare multiple schools or policies.
- Source: Parent group advocating against strict bans.
- Audience: Other parents and school board members.
- Missing context: Does not address classroom distraction or how many emergencies
  actually require direct student phone access.
- Bias: Strongly favors family access and safety concerns.

The Pattern Across Sources:
All three sources agree that phones create a real school decision, not just a
student preference. They disagree about what problem matters most. The AI gives a
balanced overview but no evidence. The news article emphasizes classroom focus and
adult management. The parent advocacy post emphasizes safety and family contact.
No single source gives a complete answer. A careful citizen would need the actual
school policy, student and parent input, teacher experience, discipline or learning
data, and exceptions for medical or safety needs.
```

This sample is **Proficient** because it applies all six questions to three source types and includes a real synthesis paragraph. It becomes **Exemplary** if the student cites the exact article/post, uses more specific evidence, and identifies what data would resolve the disagreement.

Additional Scored Exemplars

Use these shorter calibration examples when scoring varied student work. They are intentionally not perfect; they show how to score real classroom artifacts.

Exemplar A — Exemplary / advanced student

Topic: whether the county should expand bus service near a middle school.

The student analyzes: an AI overview, a county transit webpage, a local news article quoting parents, and a business association statement. For each source, the student identifies the claim, evidence, source, audience, missing context, and bias. The synthesis states: "The county page proves service changes are being studied, but it does not show student-specific safety data. The news article shows parent concern, but mostly interviews families who want expansion. The business group focuses on traffic delay, not student access. The AI answer helped me identify search terms but was not evidence."

Score: Exemplary. The student uses four source types, distinguishes evidence from advocacy, names bias in every source, and explains what no single source provides. Award Exemplary even if the writing has minor grammar errors.

Exemplar B — ELL-adapted Proficient

Topic: plastic waste at a local park.

The student uses sentence frames: "The claim is…" "The source wants…" "The missing context is…" The student analyzes an AI response, a city parks page, and an environmental group page. Some sentences are short: "The city page says cleanup days happen. It does not say how much trash is found. Bias: city wants people to think they are working on it." The synthesis says the city and advocacy source agree litter is real but disagree on urgency.

Score: Proficient. The artifact meets the analytical target even with simplified language. Do not lower the score for sentence-frame use, accent marks, or grammar that does not block meaning.

Exemplar C — IEP-modified Developing moving toward Proficient

Topic: phones in school.

The student completes a two-source version using a teacher-provided template. The AI source and a school policy excerpt are analyzed. Four of the six questions are answered for both sources; missing context is answered only once. The synthesis says: "The AI is general. The school policy is official but does not explain why the rule exists."

Score: Developing, or Proficient if the student's IEP modification reduced the required source count and the missing-context answer was orally explained. Record the accommodation. The key scoring distinction is whether the student performed the six-question reasoning, not whether the artifact matches the full template length.

Extension Options

Extension 1: Compare across time. Have students take a single topic and analyze how a single news outlet has covered it over time. Has the framing shifted? What has been added or dropped from the coverage?

Extension 2: Investigative source-tracing. Take an AI claim or social media post; trace its assertions back through the chain of sources to the original. How many hops does it take? Where does the claim get distorted along the way?

Extension 3: Read across the political spectrum. Have students analyze the same topic across one left-leaning, one right-leaning, and one center-positioned news source. What do all three agree on? What does each emphasize that the others ignore? Particularly suitable for advanced civics students.

Extension 4: Foundation Edition Chapter 8 reading. The research methodology in Chapter 8 of *The Open Source Student* (Foundation Edition) extends today's analytical work into formal research practice.

No-Install Adaptation Notes

If you do not have live AI access for the demonstration, use prepared materials.

Sample topic and prepared source set: Sea level rise impacts in Florida

For the Activity 1 demonstration, you can use this prepared source set on the topic of sea level rise impacts in Florida. The sources illustrate the six-question framework cleanly.

Source 1 — AI response (prepared): *"Sea level rise is a major concern for Florida. Scientists project that sea levels along Florida's coast will rise by approximately 1 to 4 feet by 2100, with significant impacts on coastal communities, infrastructure, and ecosystems. Cities like Miami and Tampa face particular vulnerability due to their low elevation and porous limestone bedrock. Adaptation efforts include seawalls, elevated infrastructure, and managed retreat from the most vulnerable areas."*

Six-question demonstration: - Claim: sea level will rise 1-4 feet by 2100; specific FL cities are vulnerable; specific adaptation strategies exist - Evidence: vague "scientists project"; no specific study cited - Source: AI; we don't know which scientific projections it drew on - Audience: general reader, no specific group - Missing context: range of projections (1-4 feet is wide); time scale (this century is long); regional variation; political context of adaptation funding - Bias: framed as factual / neutral; AI may reflect mainstream scientific consensus, which is itself a kind of position

Source 2 — news article excerpt (prepared): *"South Florida is increasingly experiencing 'sunny day flooding' — flooding during high tides that occurs without rainfall. The phenomenon, driven by sea level rise, is now occurring on more than 60 days per year in some Miami-Dade neighborhoods, up from fewer than 10 days per year a generation ago. Local officials are debating funding for elevated streets, raised sea walls, and other infrastructure adaptations."* [Source attribution: a Florida-based news outlet]

Six-question demonstration: - Claim: sunny day flooding has increased substantially; specific number (60+ days vs. <10); local officials debating funding - Evidence: specific numbers given; "a generation ago" is vague but quantitative claim is stronger than Source 1 - Source: news outlet (named); reporter accountability for accuracy - Audience: South Florida residents and policy followers - Missing context: which specific data source for the 60-day claim; how funding debates are progressing; alternative perspectives on solutions - Bias: implicit assumption that sea level rise is the driver (which scientific consensus supports but isn't argued in the excerpt); no skeptical perspective offered

Source 3 — advocacy organization position (prepared): *"Florida coastal residents face an existential threat from accelerating sea level rise. Without immediate, large-scale federal investment in adaptation infrastructure, hundreds of thousands of homes will become uninsurable within two decades, triggering a coastal real estate collapse. We call on Florida's congressional delegation to support legislation funding $50 billion in adaptation investments."* [Source attribution: a hypothetical environmental advocacy organization]

Six-question demonstration: - Claim: existential threat; large federal investment needed; specific consequence (uninsurable homes) and timeline (two decades) - Evidence: claims not specifically sourced in the excerpt; would need to follow up - Source: advocacy organization; clear stake in the outcome (federal funding for adaptation) - Audience: congressional members and Florida voters who could pressure them - Missing context: alternative views on adaptation strategy; cost-benefit analysis of federal vs. state vs. local funding; uncertainty in projections - Bias: clearly advocating a specific policy position; not pretending to be neutral; this is honest framing of a clear stake

Source 4 — social media post (prepared): *"My neighbor's house in Miami flooded again last week. Third time this year. The water level is just getting higher. People keep building new condos like nothing's wrong. We need to wake up."* [Source attribution: a hypothetical local resident's social media post]

Six-question demonstration: - Claim: personal experience of repeated flooding; broader assertion that development continues despite the problem; call for action - Evidence: personal observation only; not verifiable from the post - Source: individual; emotional / experiential perspective - Audience: the poster's followers; broader public if shared - Missing context: exact location; whether flooding events are connected to sea level rise vs. drainage issues vs. specific weather; broader development context - Bias: clear personal frustration; legitimate experience but limited to one observer's view

Synthesis demonstration: Across these four sources, the agreement: sea level rise is happening and affecting Florida. The disagreements: how urgent it is, what to do about it, who should pay, and how confident the projections are. What no single source provides: a complete picture combining the personal lived experience (Source 4), the local empirical data (Source 2), the policy debate (Source 3), and the scientific projection (Source 1) into a coherent civic understanding. A careful citizen would need all four perspectives to form a defensible view of the issue.

Common Misconceptions

"Mainstream sources are unbiased; alternative sources are biased." All sources are biased. Mainstream sources have establishment bias; alternative sources have anti-establishment bias. The skill is identifying bias regardless of where it sits.

"If I disagree with a source, it's biased; if I agree with it, it's true." This is the most important misconception to surface. Apply the framework to sources the student likes; bias is often most visible there.

"Identifying bias means I'm being political." No — identifying bias is intellectual honesty. Refusing to identify bias in sources you trust is the political move.

"All sources are equally biased; we can't trust anything." No — bias varies in degree, transparency, and impact. A clearly-labeled opinion column is different from a news article presenting opinion as fact. The skill is calibrated reading, not blanket skepticism.

“Missing context means the source lied.” Not necessarily. Sources have to leave things out — they can’t include everything. The question is whether the missing context distorts the meaning. Some omissions matter; others don’t. Judgment is required.

Cross-Curricular Connections

Social Studies (Civics): SS.7.CG.2.9 (analyzing media for bias, symbolism, propaganda — primary), SS.7.CG.2.8 (media/individuals/interest groups influencing government — supporting), and grade 8 civic-participation extensions (SS.8.CG.2.4) are the closest fits. The Information Integrity Lab Mini is essentially the methodology for the bias-and-influence analysis civics standards expect at grade 7.

ELA (Reading and Communication): the analytical reading work directly serves B.E.S.T. ELA Reading Informational Text standards (ELA.7.R.2.4 Argument, ELA.7.R.3.3 Comparative, ELA.7.R.3.4 Rhetoric). The synthesis paragraph serves Communication — Argument (ELA.7.C.1.3) standards. Note: B.E.S.T. ELA uses “C” (Communicating) as the strand prefix for what other frameworks call “Writing”; the substance is the same.

Social Studies (History): the same six-question framework applies to historical sources. Coordinate with the social studies teacher to reinforce.

Science: scientific claims (especially in policy contexts) require this kind of analysis to navigate well. Coordinate with the science teacher about how to evaluate scientific claims in news coverage.

Parent Communication Notes

Most parents will value this lesson but it can occasionally surface concerns:

A parent worries the lesson is “teaching kids to distrust the news.” Honest response: the lesson teaches calibrated reading, not blanket distrust. Students who can evaluate sources are better positioned to identify reliable journalism than students who can’t. Reading critically is not the same as cynicism.

A parent worries the lesson is “political.” Honest response: the lesson is non-partisan in its examples and topics. The skill is the goal; specific political conclusions are not. The framework applies to sources across the political spectrum.

A student becomes upset when finding bias in a source they trust. This is a learning opportunity. Acknowledge that this is the hardest part of the work. The student is doing it correctly.

Closing the lesson well

Lesson 3 is the curriculum’s most ambitious analytical lesson. Lesson 4 shifts focus to the personal — privacy, ethics, academic integrity. The transition matters.

End-of-lesson language: “You have just done some of the hardest work in this curriculum. You looked at multiple sources critically, named bias including in places it was uncomfortable to name, and started building the analytical habits that distinguish a thoughtful citizen from a passive consumer of information. Tomorrow we turn the analysis inward. We look at how YOU use AI — privacy, ethics, integrity — and what it means to be the kind of person who uses these tools well.”

Lesson 4: Privacy, Ethics, and Academic Integrity

Teacher Edition

Lesson at a Glance

Lesson length: 45 minutes core instruction + 20 minutes student independent work

Learning objectives. By the end of this lesson, students will be able to: 1. Identify categories of personal information that should not be entered into commercial AI services 2. Distinguish appropriate from inappropriate AI use in school work 3. Articulate the capability-vs-credential trap and explain why it matters 4. Produce an AI Use Disclosure that demonstrates honest documentation of AI use 5. Connect personal ethics around AI to broader civic responsibility

Florida standards addressed: - **FL B.E.S.T. CS (grades 6–8):** Personal Health and Safety — SC.7.HS.1.4 (categorize potential dangers), SC.7.HS.1.6 (risks/benefits of Internet access), SC.7.HS.1.7 (safe practices), SC.8.HS.1.2 (procedures to protect personal information, primary), SC.8.HS.3.2 (digital media influence on behavior); Cyber Security — SC.7.CS.1.1 (data in three states), SC.7.CS.3.3 (data vulnerabilities), SC.8.CS.2.1 (network/IoT privacy), SC.8.CS.2.3 (data permanency); Technological Impact — SC.7.TI.2.1 (legal/ethical behaviors), **SC.7.TI.2.2 (responsible use of communication media — strong primary fit)**, SC.7.TI.2.4 (ethical use of collected data); **Computing Components — SC.8.CO.3.1 (desktop applications vs. complementary online subscription version — strong primary fit for local-AI vs. cloud-AI privacy framing)** - **Florida Civics:** SS.8.CG.2.2 (responsibilities of citizens at local/state/national levels) and SS.8.CG.2.3 (civic virtue) — grade 8 extensions taught within M/J U.S. History #2100010 - **B.E.S.T. ELA:** Communication — Argument (ELA.7.C.1.3 — note B.E.S.T. uses C.1 prefix, not W.1); Vocabulary (ELA.7.V.1.1); K-12 ELA Expectations — Collaborative Discussion (ELA.K12.EE.4.1) and Voice & Tone (ELA.K12.EE.6.1)

(See the Standards Alignment Matrix for benchmark-level alignment.)

Materials needed: - Teacher device with browser/projection - Pre-prepared scenarios for Activity 1 discussion (provided in this Teacher Edition) - Student copies of Lesson 4 Student Edition - AI Use Disclosure template (in Student Edition) - Optional: copies of school's existing AI use policy or academic integrity policy

Google Classroom quick access: Google Classroom Upload Pack — setup guide, copy/paste assignment posts, upload manifest, and Bellringer / Warm-Up pacing options.

Prerequisites: Lessons 1, 2, and 3 complete. The verification (Lesson 2) and analysis (Lesson 3) work informs the ethical reasoning in this lesson.

Vocabulary introduced: Privacy, Personal information, Disclosure, Plagiarism, Academic integrity, Capability

Portfolio artifact produced: AI Use Disclosure (Lesson 4 contribution)

Assessment: AI Use Disclosure evaluated with rubric below

Background for the Teacher

This is the curriculum's most important lesson for protecting students from the most common ways AI use damages them — both in the immediate term (privacy violations, academic integrity violations) and in the long term (the capability-vs-credential trap that hollows out their actual learning).

Three framings carry the lesson:

Privacy is structural. The lesson is not about telling students "don't trust AI companies" — it's about recognizing that commercial AI services are not designed to keep their user data private and adapting your behavior accordingly. The simple test ("would I be comfortable if this were posted publicly with my name?") is the working heuristic students can apply in real time without needing to read every privacy policy.

The capability-vs-credential trap is the deepest argument in the curriculum. Students who internalize this concept use AI well throughout their education and beyond. Students who don't internalize it use AI to skip the cognitive work that school is supposed to develop, and pay the price years later when capability gaps surface. The lesson treats this with care — not moralizing, just honest about consequences.

Disclosure is the move that keeps students safe. Hidden AI use is high-risk academically; honest disclosure of permitted use is low-risk. The disclosure habit, built early, is what makes ongoing AI use in school work sustainable. This is also the practice that builds character — students who learn to handle their methods honestly become adults who handle their methods honestly.

The lesson is intentionally non-judgmental about students who have used AI inappropriately in the past. Most students will have. The goal is to build the habit going forward, not to punish past mistakes. The "what if you made a mistake" section in the student reading exists for this reason.

Lesson Procedure

Total: 45 minutes core instruction + 20 minutes student independent work.

Timing note for first delivery: Times are targets, not requirements. First-time delivery often runs 5–10 minutes long because students want to discuss edge cases around academic integrity. If pacing gets tight, use three scenarios instead of five and assign the disclosure reflection as independent completion work.

Bellringer / Warm-Up Options (choose one)

Use one option at the start of class. If pacing is tight, use Option A and treat it as the first three minutes of the Opening rather than an added activity. Use this rating system: OK to share = generally safe to type into an approved AI tool; Use caution = maybe safe only after removing personal details or asking your teacher;

Do not share = too personal, private, or risky to type into an AI tool. Use this item list: (1) a public article link your teacher gave you, (2) your full name, school, and daily schedule, (3) a paragraph you wrote for class, (4) a private story a friend told you, and (5) a question about a vocabulary word.

Option A — Fast Start (3 minutes): Post the rating system and item list above. Students choose one item, mark it OK to share, Use caution, or Do not share, then explain why in one sentence.

Option B — Standard Warm-Up (5 minutes): Post the rating system and all five items above. Students classify every item as OK to share, Use caution, or Do not share. Ask for quick share-outs that include one safe item and one unsafe item.

Option C — Extended Launch (6-7 minutes): Post the rating system and all five items above. Students classify every item, then write one rule that could help a student decide what not to type into an AI tool. Use their rules to introduce the lesson's privacy heuristic: if you would not want it posted publicly with your name, do not put it into an AI tool.

Opening (5 minutes)

Recap from Lessons 1–3: "You've learned that AI hallucinates, that you should verify, and that you should read all sources critically. Today we turn the focus inward: how do YOU use AI in ways that make you stronger, not weaker?"

Bridge: "By the end of class, you'll have produced an AI Use Disclosure — a documented record of how you've been using AI in your school work. This is the practice that makes your AI use honest and sustainable."

State lesson goal as written in Student Edition.

Direct Instruction (15 minutes)

Walk students through the Student Reading. Pace carefully — the content is dense and consequential.

Key sections to emphasize:

"Privacy: What Not to Share" — read the categories aloud. Pause: "Has anyone here typed something into AI that, looking back, you wish you hadn't?" Don't require sharing; the question itself is the teaching moment. The simple test ("would I be comfortable if this were posted publicly with my name?") is the heuristic students should leave with.

"When You Use Local AI, Privacy Works Differently" — short paragraph, important distinction. Local AI has different (better) privacy properties than commercial AI. This explains why the curriculum's Foundation Edition emphasizes local AI installation.

"Ethics: Using AI Honestly in School Work" — the central content of the lesson. Read the two-students example carefully. Pause and ask: "Which student is in better shape five years from now? Why?" The answer should be Student B, but get students to articulate why — because Student B has both credential and capability, while Student A has only credential.

“The capability-vs-credential trap” — the most important conceptual content in the lesson. Pause after reading. Ask: “Have you seen this happen to anyone you know? Have you done it yourself, even a little?” This is the moment many students will privately recognize themselves; the recognition is part of the learning.

“When AI Use Is Clearly Fine, Clearly Not Fine, and Gray” — read the lists. Acknowledge the gray areas honestly. The “ask the teacher” guidance is critical; students who don’t know the answer should ask, not guess.

“Disclosure: The Honest Move” — read the disclosure example. This is the format students will use in Activity 2.

“What If You Made a Mistake?” — important paragraph for any student who has used AI inappropriately in the past. The honest move is to come forward; the consequences of disclosure are almost always smaller than consequences of being discovered.

Activity 1: Discussion of AI Use Scenarios (10 minutes)

Present 4–5 scenarios for class discussion. For each: was this use appropriate? Why or why not? If gray, what would make it clearly appropriate or not?

Scenario 1: A student has trouble understanding a chapter on cell biology. They ask AI to explain the chapter’s main concepts in simpler terms. They then re-read the chapter and feel they understand it. They write their own summary for class.

[Discussion prompts for educator: This is a clearly appropriate use. AI was used for understanding, not for producing the submission. The student did the cognitive work. No disclosure required for most assignments because the use was support for understanding, not production.]

Scenario 2: A student has an essay due tomorrow. They give the prompt to AI, take the AI’s response, change a few sentences to sound like their voice, and submit it.

[Discussion prompts: Clearly inappropriate. The submitted work is essentially AI’s, not the student’s. The “change a few sentences” doesn’t make it the student’s work. This is the textbook capability-vs-credential failure. If the student had disclosed honestly, they would still be in trouble; the violation is the substitution itself, not just the failure to disclose.]

Scenario 3: A student is brainstorming a research topic. They ask AI for 10 possible angles on the broad topic. They pick one of the suggested angles, then go research and write the paper themselves.

[Discussion prompts: Generally appropriate, especially with disclosure. Brainstorming with AI is comparable to brainstorming with a friend or sibling — the ideas don’t make the paper any less the student’s work. Disclosure is the responsible move when the chosen angle came from AI suggestion.]

Scenario 4: A student is working on a math problem set. They paste the problems into AI to get the answers. They then write the answers on the homework worksheet without working through the problems themselves.

[Discussion prompts: Clearly inappropriate. The cognitive work the homework was supposed to develop has been skipped entirely. The student has the credential of completed homework and zero capability gain.]

Scenario 5: A student is writing a short story for English class. They use AI to help generate names for their characters and descriptions of fictional places. They write the story themselves.

[Discussion prompts: Generally fine, especially with disclosure. Creative-task assistance is comparable to using a baby names website or a thesaurus. The student is doing the storytelling; AI helped with peripheral details. Disclosure is courteous.]

Scenario 6: A student is confused by a math procedure. They ask AI to explain the first step in simpler language, close the AI tool, then solve the assigned problems by showing their own work.

[Discussion prompts: Appropriate if teacher policy allows AI explanation support. This builds capability because the student still practices the target skill. The same scenario becomes inappropriate if the student copies AI's worked answers without doing the reasoning.]

Scenario 7: A student asks AI to generate an essay outline, then fills in each paragraph with AI-generated topic sentences and lightly changes the wording before submitting.

[Discussion prompts: Usually inappropriate unless the assignment explicitly allows that level of AI drafting. The student may be doing some surface editing, but AI is controlling the structure and much of the substance. This is a credential risk because the submitted work may overstate the student's actual writing and planning capability.]

Scenario 8: A student uses AI to create practice questions before a civics quiz, answers the practice questions without AI, checks the textbook when unsure, and studies the mistakes.

[Discussion prompts: Appropriate and strongly capability-building. AI generated practice, but the student did the retrieval, checking, and learning work.]

The discussion takes about 10 minutes if you keep moving. Some scenarios will produce immediate consensus; others will produce real debate. The debate is the teaching moment.

Activity 2: Student AI Use Disclosure (15 minutes core, plus completion as homework)

Direct students to Activity 2. Each student produces their own AI Use Disclosure based on their actual AI use across this curriculum (and other school work where they've used AI).

Setup (3 minutes): Distribute (or direct to) the AI Use Disclosure template. Walk through the template structure.

Student work (10 minutes): - Students reflect honestly on their AI use - They complete the template - They write the short reflection

Honesty norms during Activity 2: - Be clear with students about how honest disclosure will be handled in your specific setting. Disclosure should be evaluated according to your school's academic integrity policy — not under a classroom-level amnesty unless the school has explicitly approved one. The lesson's substantive point (honest disclosure is the right move and produces better long-term outcomes than concealment) does not require promising any specific consequence outcome - The grading criterion for the disclosure activity itself is honesty of process, not whether students used AI a lot or a little - Students may keep their disclosure private if they prefer (submit to teacher only) or share with the class

Wrap-up (2 minutes): - One or two students share what was hardest about writing their disclosure - Acknowledge: this is the kind of practice that builds real trust between students and teachers

Closing (5 minutes)

Bring class back. Acknowledge that today's content was emotionally heavier than other lessons — privacy, ethics, integrity, the capability trap. These are real stakes and students felt that.

Final framing: "You now have the foundation: you know what AI is and isn't, you know how to verify, you know how to read across sources, and you know how to use AI ethically. Tomorrow we put it all together in a project. You'll pick a real issue in your community or state, research it using the skills you've built, and produce something that demonstrates everything you've learned. This is the capstone of the pilot."

Direct students to Reflection Questions and Self-Check.

Student Reading — Reproduced with Teacher Annotations

> **This lesson is about using AI in ways that strengthen you rather than hollow you out.** [TEACHER NOTE: The lesson's central framing. Underline.]

> **Don't share information that identifies you specifically.** [TEACHER NOTE: The most concrete privacy guidance. Read the categories aloud. Many students have already typed personal info into AI without thinking; this is the moment they may realize the implications.]

> **The simple test:** [TEACHER NOTE: The heuristic students should remember. Repeat it during direct instruction.]

> **The ethics question for AI use in school is not "should I use AI?"** [TEACHER NOTE: This reframing matters. The question is "how do I use it well?"]

> **This is the capability vs. credential trap.** [TEACHER NOTE: The single most important conceptual content in the lesson. If students remember nothing else, they should remember this trap and how to avoid it.]

> **When AI Use Is Clearly Fine, Clearly Not Fine, and Gray** [TEACHER NOTE: The lists are pragmatic. The “ask the teacher” guidance for gray areas is the most important: when in doubt, ask.]

> **Disclosure is the move that keeps you on the right side of academic integrity rules.** [TEACHER NOTE: This is the practice that makes ongoing AI use sustainable. Hidden use is high-risk; honest disclosure is low-risk.]

> **What If You Made a Mistake?** [TEACHER NOTE: This paragraph exists for students who recognize themselves in the inappropriate-use scenarios. The honest move is to come forward. Most teachers respond much better to disclosure than to discovery.]

Vocabulary Teaching Strategies

Most terms in today’s lesson will be familiar from social studies or general usage. The two terms requiring focused teaching:

“Capability vs. credential” — the conceptual heart of the lesson. Use the two-students example. Make sure students can distinguish: a credential is documentation that you completed something; capability is the actual skill the documentation was supposed to represent. They are not the same thing; AI use can produce one without the other.

“Disclosure” — students may know the word in other contexts (legal disclosure, financial disclosure). In the AI use context, it means clearly explaining where AI helped with your work. The disclosure example in the reading is the format.

For ELL students: “privacidad,” “información personal,” “divulgación,” “plagio,” “integridad académica” all map cleanly. “Capability” maps to “capacidad.”

Differentiation Notes

Below grade level readers

- Read aloud the privacy categories
- Use the two-students story as an oral discussion before students read it themselves
- Simplify the AI Use Disclosure to a checkbox version

Above grade level readers

- Foundation Edition Chapter 14 is the deeper version of today's content
- Extension: have them write a longer reflection on a time they (or someone they know) faced the capability-vs-credential trap and what they would do differently

English Language Learners

- All vocabulary has cognates; flag for ELL students
- Pair for the discussion activity
- Allow the AI Use Disclosure to be completed in the student's stronger language with key terms in English

IEP / 504 — ADHD / Executive Function

- Use the lists in the reading as visible checklists
- Break Activity 2 into smaller segments

IEP / 504 — Reading / Language Processing

- Audio of reading
- Allow oral disclosure

Autism Spectrum

- The systematic ethical framework may be a strength
- Allow private completion of disclosure if discussion is overwhelming

Anxiety / Emotional Regulation

- Be clear with students about how honest disclosure will be handled in your specific school setting (per your school's academic integrity policy); reduce anxiety by surfacing the actual policy rather than improvising a classroom-level promise the school has not approved
- Disclosure can be private (to teacher only) if student prefers

Limited Technology Access

- The lesson works without any AI access
- AI Use Disclosure can be on paper

Gifted Students

- Have them research how journalism organizations or publishing houses handle AI disclosure
- Have them propose an institutional AI use policy for their school or microschool

---

Assessment Rubric: AI Use Disclosure

Evaluate primarily on honesty and completeness, not on whether the student used AI a lot or a little.

Criterion	Developing	Proficient	Exemplary
Identifies AI services used	Vague or absent	Names specific AI services used	Names services and notes when each was used
Documents specific uses	Generic “I used AI for help”	Specific tasks for which AI was used	Specific tasks with description of how AI’s output was used (verbatim, as starting point, for feedback only)
Confirms own work	Confirmation is rote or unclear	Clear confirmation that substantive work is the student’s own	Confirmation distinguishes between AI-supported and AI-substituted work; honest about boundary
Reflection: where AI helped learning	Generic (“It helped me”)	Specific example of AI supporting learning	Specific example with reflection on why this use built capability rather than substituted for it
Reflection: where student avoided AI	Missing or generic	Specific example of choosing not to use AI to develop the skill	Reflection includes why this choice was worth making
Reflection: what student would do differently	Vague or absent	Specific change the student plans to make	Change is grounded in the capability-vs-credential framework
Honesty	Disclosure feels rehearsed or sanitized	Disclosure reads as genuine	Student includes uncertainty or past mistakes; demonstrates real reflection

Answer Key / Scoring Notes

Lesson 4 is graded on judgment, honesty, and disclosure quality. Students do not need to have used AI heavily to complete the artifact well.

Warm-Up classification key: - Public article link from teacher: **OK to share** with an approved AI tool, assuming the article is public and the teacher allows AI use. - Full name, school, and daily schedule: **Do not share**. It identifies the student and creates safety/privacy risk. - Paragraph the student wrote for class: **Use caution**. It may be okay if assignment policy allows it and personal information is removed; it should not be used to outsource the work. - Private story a friend told the student: **Do not share**. It is someone else’s private information. - Question about a vocabulary word: **OK to share** with an approved AI tool; this is a learning-support use.

Activity 1 scenario key: 1. Cell biology explanation: **appropriate**. AI supports understanding; student still does the reading and writes their own summary. 2. Essay generated by AI and lightly reworded: **inappropriate**. AI substitutes for the student’s work; rephrasing does not fix authorship. 3. Brainstorming research angles: **generally appropriate with disclosure** if the student researches and writes independently. 4. Math answers copied from AI: **inappropriate**. The assignment’s purpose is the student’s mathematical practice. 5. Character names and fictional place descriptions: **generally appropriate with disclosure** if the student writes the story and uses AI for peripheral support.

Self-Check answer key: 1. Personal information not to type into commercial AI includes full name, address, phone number, school schedule, student ID, passwords, private family information, medical information, financial information, private stories about friends, photos of private people, or anything the student would not want posted publicly with their name. 2. A credential is proof or appearance of completing work; capability is the actual skill or understanding the work is supposed to build. 3. Simple test: if you would not

be comfortable with the information posted publicly with your name, do not put it into an AI tool. 4. An AI Use Disclosure names the AI tool used, what it was used for, what parts of the work are the student's own, how the student verified or revised the output, and what boundaries were followed. 5. The right move is honest disclosure to the teacher or appropriate adult, following the school's academic integrity policy. The student should explain what happened, what AI did, what the student did, and what they will do differently.

Sample Proficient Artifact (for educator reference)

```
AI USE DISCLOSURE

Student: A.T.
Assignment / Project: Civic Tech Mini-Project planning notes
Date: May 14, 2026

AI service used:
ChatGPT, during class demonstration and once at home with parent permission.

What I used AI for:
I asked AI to help me list possible local issues connected to water quality and
school transportation. I also asked it to explain the difference between a city
department and a county department.

What I did not use AI for:
I did not ask AI to write my final paragraph or make my position for me. I did
not paste my full name, school schedule, address, or private information into AI.

How I checked the AI output:
The AI suggested checking my county environmental services page. I verified the
agency name on the county website. One AI suggestion about a state office was too
general, so I did not use it in my brief.

What part is my own work:
The source choices, notes, final position, and writing are mine. AI helped me
brainstorm and understand terms, but I made the decisions and wrote the final
artifact.

What I learned:
AI is helpful when I am stuck at the beginning, but it can make things sound more
certain than they are. Next time I will ask AI for questions to investigate, not
for answers to copy.
```

This sample is **Proficient** because it names the service, describes specific uses, confirms student authorship, and explains verification. It becomes **Exemplary** if the student names exact prompts or timestamps, includes a clearer privacy boundary, and connects the reflection explicitly to the capability-vs-credential trap.

Additional Scored Exemplars

Use these examples to calibrate disclosures that are honest but vary in detail, language level, or accommodation level.

Exemplar A — Exemplary

The student names the tool, date, purpose, and boundary: “I used Gemini on May 13 to ask for three possible questions about school bus routes. I copied none of its wording. I used one suggestion to search the county transportation page. I did not enter my name, school schedule, address, or my friend’s story about missing the bus. The final question and paragraph are mine.” The reflection says AI helped with “starting questions” but not with “the thinking that proves I understand the issue.”

Score: Exemplary. Specific tool/date/purpose, clear privacy boundary, clear authorship boundary, and explicit capability-vs-credential reflection.

Exemplar B — ELL-adapted Proficient

The student writes: “I used AI for vocabulary words: agency and jurisdiction. I asked in English and Spanish. I did not put my full name or address. I wrote my own answer. AI helped me understand, but I did the project.” The student adds a short oral explanation to the teacher: “I did not use AI to write my position.”

Score: Proficient. The disclosure is short but complete. Accept bilingual wording or oral clarification when the student demonstrates the authorship and privacy distinction.

Exemplar C — IEP-modified Developing moving toward Proficient

The student completes a checkbox disclosure: - AI used: yes - Used for: understand words / brainstorm questions - Did not use for: final answer - Private information shared: no - My own work: final paragraph and source choices

The written reflection is one sentence: “Next time I will ask before using AI.”

Score: Developing on the standard rubric because reflection is thin; Proficient may be appropriate if the student’s accommodation permits reduced written output and the teacher records an oral reflection explaining why AI support did not replace the student’s work.

Exemplar D — Beginning / needs reteach

The student writes only: “I used AI but I changed it so it is mine.”

Score: Beginning. This does not identify the tool, task, boundary, verification, or authorship distinction. Reteach the capability-vs-credential trap and require a revised disclosure.

Extension Options

Extension 1: Have students review their school’s existing academic integrity policy and draft a paragraph that would update it for the AI era.

Extension 2: Have students write a short letter to a younger student who is just starting to use AI, advising them on privacy and ethics.

Extension 3: Foundation Edition Chapter 14 reading. Significantly more depth on the same content.

Extension 4: Research how professional fields handle AI disclosure (journalism, academic publishing, legal practice). Compare their norms to what students should do in school work.

No-Install Adaptation Notes

This lesson does not require live AI access for any activity. The discussion scenarios in Activity 1 are pre-prepared. The AI Use Disclosure in Activity 2 is reflective writing.

The only AI-related need is that students should have actually used AI at some point in their school work to have something to disclose. If a class includes students who have not yet used AI, those students can complete the disclosure for the AI use within this curriculum (the AI demonstrations they observed, any AI interactions during Lessons 1–3 activities) and reflect on how they would handle disclosure if they began using AI more.

Common Misconceptions

"Privacy doesn't matter because I don't have anything to hide." This framing misunderstands privacy. Privacy isn't about hiding wrongdoing; it's about controlling what information about you is available to whom. Even people with nothing to hide have legitimate reasons to keep their information out of databases they don't control.

"If I rephrase what AI gave me, it's my work." Rephrasing doesn't transform AI's work into your work. The thinking and the substance remain AI's. The capability-vs-credential trap operates regardless of how you transform the surface text.

"Using AI is cheating." Not necessarily. The lesson distinguishes between AI use that builds capability (legitimate) and AI use that substitutes for capability (cheating). Using AI to understand a concept is fine; using AI to skip the writing is not.

"My teacher won't notice if I don't disclose." Teachers are increasingly able to detect AI-generated work, and they are also able to test whether students can defend their submitted work. The "won't notice" assumption is increasingly wrong; honest disclosure remains the safer practice regardless.

"If AI is okay sometimes, it's okay all the time." No — appropriateness varies by assignment, context, and what the assignment is meant to develop. The "ask the teacher" guidance for gray areas is essential.

Cross-Curricular Connections

ELA (Writing): academic integrity in writing is foundational across all subjects. The AI disclosure practice transfers to any writing assignment.

Social Studies (Civics): civic responsibility and personal integrity are linked. The lesson's broader framing — that personal ethics around AI shape civic responsibility — connects directly to civics standards.

Health: privacy and digital wellness are part of middle school health curriculum. The lesson's privacy content reinforces health curriculum themes.

All academic subjects: the capability-vs-credential framing applies to every subject. Coordinate with teachers across subjects to reinforce the framework.

Parent Communication Notes

This lesson can produce parent questions about AI use at home and at school.

A parent asks for the lesson's framework so they can apply it at home. Share the categories of personal information not to share with AI, the simple test, and the capability-vs-credential framing. Most parents value all three.

A parent expresses concern about academic integrity. Direct them to today's lesson content and the Student AI Use Agreement. The curriculum's framework is conservative and aligned with traditional academic integrity principles.

A parent worries their student has used AI inappropriately. Reinforce the "what if you made a mistake" framing — honest disclosure is the path forward. Acknowledge that most students will have used AI in some inappropriate way at some point; the goal is the habit going forward, not punishment for past mistakes.

Closing the lesson well

End-of-lesson language: "Today's content was heavier than rest of the week. Privacy, ethics, integrity — these are real stakes for who you are becoming. The students who handle these well in school become adults who can be trusted with much more. The students who don't tend to limit their own future possibilities. This is up to you, every time you sit down to use AI. Tomorrow we put everything together in the Civic Tech Mini-Project — your chance to use these skills for something that matters in your community."

Lesson 5: Civic Tech Mini-Project

Teacher Edition

Lesson at a Glance

Lesson length: 60–75 minutes core instruction recommended (longer than other lessons because it includes a project) + 30+ minutes student independent work for project completion

Learning objectives. By the end of this lesson, students will be able to: 1. Identify a real local or state civic issue suitable for short-form research 2. Frame an issue as a specific researchable question 3. Identify the responsible government agency or official for a chosen issue 4. Apply verification, source analysis, and ethical AI use practices to a real research project 5. Produce a Civic Issue Brief that demonstrates evidence-based civic engagement 6. Disclose AI use honestly within a substantive piece of work

Florida standards addressed (capstone — touches all major frameworks): - **FL B.E.S.T. CS (grades 6–8):** Technological Impact — engaging public officials (SC.8.TI.2.2, near-perfect fit) and responsible communication-media use (SC.7.TI.2.2); Emerging Technologies (SC.7.ET.2.1 — AI as research/drafting partner); Communication and Collaboration (SC.7.CC.2.2 — analyze own ideas with research-based information to create unique digital artifact) - **Florida Civics: SS.7.CG.2.10 (PRIMARY — near-verbatim benchmark match: "process for citizens to address state or local problems by researching public policy alternatives, identifying appropriate government agencies, and determining a course of action")**, SS.7.CG.2.9 (continued — analyzing media for bias), SS.7.CG.2.8 (continued — media/groups influencing government), SS.8.CG.2.2 / SS.8.CG.2.6 (grade 8 active-participation extensions) - **B.E.S.T. ELA:** Communication — Argument (ELA.7.C.1.3), Research (ELA.7.C.4.1), Oral (ELA.7.C.2.1), Multimedia (ELA.7.C.5.1), Tech in Communication (ELA.7.C.5.2); Reading Informational Text — Comparative (ELA.7.R.3.3); Vocabulary (ELA.7.V.1.1); K-12 ELA Expectations — Collaborative Discussion (ELA.K12.EE.4.1)

(See the Standards Alignment Matrix for benchmark-level alignment.)

Materials needed: - Teacher device with browser/projection - Class display - Internet access for student devices (recommended for live source-finding; **for no-internet settings, the educator pre-curates a Source Packet** — a small folder of printed government documents, news articles, and advocacy materials for each candidate issue, prepared the day before. The Source Packet path is the documented offline alternative for Lesson 5 and supports the kit's no-internet contingency claim) - Student copies of Lesson 5 Student Edition - Civic Issue Brief template (in Student Edition; can also be reproduced as worksheet) - Pre-prepared list of suggested local/state issues for students who need topic suggestions - For pilots in specific Florida counties: pre-research a handful of local issues that students might engage with (links to local government pages, current local news topics, etc.)

Google Classroom quick access: Google Classroom Upload Pack — setup guide, copy/paste assignment posts, upload manifest, and Bellringer / Warm-Up pacing options.

Prerequisites: Lessons 1, 2, 3, 4 complete. This is the capstone; all four prior lessons' skills are exercised here.

Vocabulary introduced: Agency, Jurisdiction, Policy, Public official, Public records

Portfolio artifact produced: Civic Issue Brief (Lesson 5 contribution and culminating artifact for the curriculum)

Assessment: Civic Issue Brief evaluated with rubric below

Background for the Teacher

Lesson 5 is the curriculum's most ambitious lesson logistically and pedagogically. It is also the lesson that demonstrates whether the previous four lessons actually built durable capability — students either can or cannot apply verification, source analysis, ethical AI use, and civic engagement skills to a real issue.

Three framings carry the lesson:

Real issues, not hypothetical ones. The lesson is built around real local and state issues because the work of civic literacy is meaningless if it is only practiced on hypotheticals. Students who research a real local zoning question, a real Florida bill, or a real environmental issue produce work that connects to civic life directly. The "real issue" requirement is binding, not advisory.

Short-form, not comprehensive. The mini-project is bounded — one lesson period, plus completion time. Students will not produce comprehensive policy analyses; they will produce defensible issue briefs that demonstrate they can do the work. The bounded scope is a feature, not a limitation. The full nine-week module (Phase 2) and full-year edition (Phase 3) will support more substantial civic technology projects.

Position-taking with evidence. The Civic Issue Brief includes a section where students state their position on what should happen. This is intentional. Civic literacy includes the capacity to take a defensible position, not just to neutrally report. The grading criterion is whether the position is defensible from the student's sources, not what the position is. This is important to establish explicitly with students who may worry that their teacher will grade them on agreement with a particular view.

The lesson also exercises the curriculum's two pillars (AI literacy and civic technology literacy) most explicitly. AI is genuinely useful as a research and drafting assistant for civic work; the lesson models this use while requiring verification, source citation, and disclosure. Students leave with a concrete example of what responsible AI-assisted civic work looks like.

The lesson is logistically heavier than others — more student-led work, more variation in what each student produces, more time required overall. Plan accordingly. If the standard one-week pacing has Friday for Lesson 5, schedule extra time if possible (60–75 minutes for core instruction instead of 45) so the project can be completed in class. If the project must extend into homework, that is workable but reduces the in-class collaborative element.

Lesson Procedure

Total: 60–75 minutes core instruction recommended + 30+ minutes student independent work.

Timing note for first delivery: Times are targets, not requirements. First-time delivery often runs 10–15 minutes long because issue selection and agency identification are unfamiliar. If pacing gets tight, provide a pre-curated issue list and source packets, then evaluate a compact but complete Civic Issue Brief rather than a longer project.

Bellringer / Warm-Up Options (choose one)

Use one option at the start of class. If pacing is tight, use Option A and treat it as the first three minutes of the Opening rather than an added activity. For this warm-up, a real civic issue means a problem or decision connected to the student's school, neighborhood, city, county, or state that people could research and try to improve. Examples: traffic near school, school phone policy, water quality, public park safety, bus routes, recycling, library hours, local flooding, or a proposed Florida law.

Option A — Fast Start (3 minutes): Post the definition and examples above. Students name one real civic issue people might disagree about or try to improve. Scan responses quickly to identify students who may need topic support.

Option B — Standard Warm-Up (5 minutes): Post the definition and examples above, then ask: "1. What real civic issue might people disagree about or try to improve? 2. Who is affected by this issue? 3. What public agency, public official, school board, city, county, or state office might be connected to it? 4. What source could you check first to learn more? Examples: a school board agenda, city or county website, Florida agency page, local news article, or official meeting minutes." Ask two students to share the public agency or official they think might connect to their issue.

Option C — Extended Launch (6-7 minutes): Post the definition and examples above, then ask: "1. What real civic issue might people disagree about or try to improve? 2. Who is affected by this issue? 3. What public agency, public official, school board, city, county, or state office might be connected to it? 4. What source could you check first to learn more? 5. What would make this issue too broad, too vague, or too hard to research today? How could you narrow it?" Use one broad example and one narrowed example to model the difference between a topic that is too big and a question students can actually research today.

Opening (5 minutes)

Recap the week: "You've built a set of skills this week — understanding what AI is, verification, source analysis, ethical use. Today we put it all together. You're going to take a real issue in your community or state and do real research on it, using everything you've learned. By the end of class, you'll have produced a Civic Issue Brief that documents what you found and what you think should happen."

Bridge: "This is the lesson where the curriculum's purpose becomes concrete. The skills you've built are not just academic exercises. They are what citizens do when they engage seriously with public life."

State the lesson goal as written in Student Edition.

Direct Instruction (10 minutes)

Walk students through the Student Reading. Lighter touch than previous lessons because students are about to do a substantial project; the reading is more orienting than instructional.

Key sections to highlight:

“Why a Civic Project, Specifically” — connect to civics standards explicitly. This is why Florida has the standards it has; today’s project is the practice.

“What Counts as a ‘Real’ Issue” — the criteria are non-negotiable. Real issue, locally or state-relevant, identifiable agency, real sources available.

“How to Approach the Mini-Project” — the five-step structure. Walk through each step briefly so students know the path.

“A Few Notes on Doing This Well” — the practical guidance. Stay specific. Cite sources. Take a defensible position. Be honest about uncertainty. Disclose AI use.

Activity 1: Issue Selection (10 minutes)

This is the moment many students will need scaffolding. Some will know exactly what they want to research; others will be stuck.

For students who have an issue in mind: check that it meets the “real issue” criteria. Is it specific? Is it local or state level? Can they identify an agency? Are sources findable? If yes, they’re ready to start research.

For students who don’t have an issue: offer the suggested topic categories from the reading. Have them pick one category, then narrow within it. “Local environment” is too broad; “water quality in [your local water body]” is workable.

For students who need more scaffolding: prepare in advance a list of 5–10 specific local issues you can offer. For Florida pilots, possibilities include (research these in advance for your specific county): - A current local government decision making news in your county - A current Florida legislative bill on a topic of likely student interest - A recent school board policy decision - A local environmental issue with active sources (news coverage, government documents)

Florida Regional Topic Starters

Use these as inspiration only; teachers should confirm local relevance before assigning.

South Florida / Miami-Dade / Broward / Palm Beach - Sunny-day flooding, drainage, and seawall planning - Transit access near schools and libraries - Heat safety for outdoor school activities - Everglades restoration and local water-management decisions

Central Florida / Orlando / Tampa Bay - Pedestrian safety near schools and commuter traffic corridors - Stormwater runoff into lakes - Affordable housing pressure near school communities - Public transit access for students and families

North Florida / Jacksonville / Gainesville / Big Bend - River and spring water quality - Hurricane shelter access and emergency communication - Rural broadband and library technology access - School bus route length and student transportation time

Panhandle - Hurricane recovery, beach access, and resilient infrastructure - Military-community traffic and housing pressures - Water quality in bays and coastal ecosystems - Public recreation access and park maintenance

Florida Keys - Sea level rise adaptation and road elevation - Marine sanctuary protection and tourism impact - Hurricane evacuation planning - Workforce housing for public-service workers

Teacher caution: Keep topics local, specific, and researchable. If a topic becomes too politically charged or too broad for one lesson, narrow it to a public agency question students can research today.

The goal of this activity is that every student leaves with a specific researchable question and a clear next step. Spend the 10 minutes on this; students who start without clear questions will produce thin briefs.

Activity 2: The Mini-Project (35–45 minutes core; project completion may extend beyond)

This is the core work of the lesson. Students follow the five-step structure from their reading.

Active facilitation during the project:

- **Circulate constantly.** This is high-variation independent work; students will need different kinds of help.
- **Push for specificity.** When students are working at too high a level of abstraction ("climate change," "education funding"), help them narrow.
- **Push for verification.** When students are about to use an AI claim without checking it, prompt: "What's your source for that?"
- **Push for honest position-taking.** When students are afraid to take a position, reassure: the grade is on whether the position is defensible from the sources, not on what the position is.
- **Push for AI use disclosure.** When students aren't documenting AI use, remind them.

Common student situations and responses:

- *"I can't find any sources."* — help them think about which agency or office handles the issue; almost everything has at least an agency website. If truly nothing exists online, suggest a different (better-documented) issue.
- *"AI told me this fact and I can't verify it."* — leave it out, or note it as "AI-claim, unverified" with a clear note about why. Don't fabricate sources to support unverified AI claims.
- *"I can't decide what position to take."* — that's fine. The brief can include "I am not yet sure what I think because the evidence is mixed in these specific ways..." Honesty about uncertainty is a defensible position.
- *"My issue turns out to be more complicated than I thought."* — that's a real finding. Document the complication.

Activity 3: Class Sharing (5–10 minutes if time allows)

If time permits, have 3–5 students share their issue and main finding. Listen for similarities (what issues your students collectively care about) and differences (the range of approaches and depths).

This activity is optional. If time is tight, it can be cut, but losing it loses some of the community-building value of the project.

Closing (5 minutes)

Bring class back together. Acknowledge the work: “You just did real civic work. Not perfect. Not comprehensive. Real. The work of taking an issue, researching it with verification and source analysis, using AI responsibly, and producing a documented brief that you can defend. That’s what citizenship looks like in 2026 and beyond.”

Reflect on the week: “Look at what you built across this week. Five skills, five artifacts, one capability you’ll use for the rest of your life.”

Direct attention to the Pilot Feedback opportunity: “Your educator is going to ask you about this curriculum — what worked, what didn’t, what was missing, what should be different. Your honest feedback shapes what comes next. Take it seriously.”

Direct students to Reflection Questions and Self-Check.

Student Reading — Reproduced with Teacher Annotations

> **This week you’ve built a set of skills that, combined, are the foundation of civic technology literacy:** [TEACHER NOTE: Recap of the week. Students often need to be reminded explicitly of what they’ve built — they don’t always notice the cumulative arc.]

> **Today you use all of those skills on a real issue.** [TEACHER NOTE: The “real issue” framing matters. Reinforce during direct instruction.]

> **Why a Civic Project, Specifically?** [TEACHER NOTE: The connection to Florida civics standards is explicit here. Make sure students understand why this is in the curriculum, not just what they’re being asked to do.]

> **What Counts as a “Real” Issue** [TEACHER NOTE: The criteria are binding. Real, local-or-state-relevant, identifiable agency, real sources. Most student topic selections will need refinement against these criteria.]

Step 1: Choose your issue and frame the question. [TEACHER NOTE: This is where most projects succeed or fail. A vague topic produces a vague brief. A specific question produces useful work.]

Step 2: Identify the responsible agency or official. [TEACHER NOTE: This is real civic literacy. Students who can locate the right agency for an issue have a foundational citizenship skill.]

Step 3: Find official sources. [TEACHER NOTE: Florida's Sunshine Law makes this particularly accessible in our state. Mention this — it is part of why civic research is more practical in Florida than in many states.]

Step 4: Use AI as a thinking partner — with verification. [TEACHER NOTE: This is where the curriculum's argument gets demonstrated. AI is genuinely useful for civic research; verification and source citation are what make that use responsible.]

Step 5: Produce your artifact. [TEACHER NOTE: Three options offered. The Civic Issue Brief is recommended for most students; the letter and presentation options work for students with specific interests.]

Stay specific. [TEACHER NOTE: The most common student failure mode is staying too abstract. Push for specificity throughout the project.]

Take a position, but defend it. [TEACHER NOTE: This may need explicit reassurance. Students sometimes worry teachers will grade based on the position; reassure that the criterion is whether the position is defensible from the sources.]

Be honest about uncertainty. [TEACHER NOTE: Uncertainty is intellectually honest. Students who can say "I don't know X" have done better work than students who pretend certainty.]

Vocabulary Teaching Strategies

Five terms today (agency, jurisdiction, policy, public official, public records). All are civics vocabulary students will see in social studies elsewhere; the lesson sharpens their meaning in a research context.

Agency: use a Florida-specific example. “The Florida Department of Education is an agency. The Florida Fish and Wildlife Conservation Commission is an agency. The Palm Beach County Sheriff’s Office is an agency. Each one handles a specific area of public responsibility.”

Jurisdiction: clarify the level question. “If your issue is about something the federal government decides, federal agencies have jurisdiction. If state, state agencies. If local, local agencies. Asking the wrong level of government about an issue is a common mistake. Knowing the right level is part of doing civic work well.”

Public records: Florida-specific framing. “Florida has a strong Sunshine Law. Most government documents in Florida are public records. This makes civic research easier in our state than in many places.”

For ELL students: “agencia,” “jurisdicción,” “política,” “funcionario público,” “registros públicos” all map cleanly.

Differentiation Notes

Below grade level readers

- Provide pre-selected issue options with relevant sources already gathered
- Allow students to use the Letter to a Public Official option (more concrete than brief)
- Pair below-level students with on-level partners for the project
- Reduce required source count to 3 (instead of 5)

Above grade level readers

- Push for a more substantive issue with more research depth
- Allow them to use the Class Presentation option (more advanced presentation work)
- Foundation Edition Chapter 16 covers civic engagement themes more deeply
- Have them produce a more complete Civic Issue Brief (full template, all sections substantive)

English Language Learners

- Allow Spanish-language sources for issues with Spanish-speaking community impact
- Pair with bilingual peers
- Provide the Civic Issue Brief template in two-column format (English and Spanish columns) for newcomer ELL students

IEP / 504 — ADHD / Executive Function

- Use the five-step structure as an explicit checklist
- Break the work into shorter visible segments with timer cues
- Provide pre-formatted Civic Issue Brief template with sections clearly demarcated
- Allow extended time for project completion if needed

IEP / 504 — Reading or Language Processing

- Allow oral presentation as the artifact (Letter or Presentation option works better than Brief for some students)
- Audio summaries of sources where possible
- Pair with on-level peer for source reading

Autism Spectrum

- The structured five-step process works well for many autistic students
- Allow private project work if class collaboration is overwhelming
- Provide explicit social expectations for the optional sharing activity
- A special-interest connection to the chosen issue can produce particularly engaged work

Anxiety / Emotional Regulation

- Reassure that the grade is on the process and the defensibility of the position, not on the position itself
- Allow private submission of the brief (not shared with class)
- Provide topic options that are less politically charged for students who would find charged topics anxiety-producing

Limited Technology Access

- Provide pre-printed source packets for the chosen issue
- Allow paper-based brief
- The Letter option works particularly well for students without device access (a letter to an official is naturally a paper artifact)

Gifted Students

- Allow them to take on a more complex issue (one with genuine policy debate, multiple stakeholders, contested evidence)
- Encourage longer-form briefs or actual letters they could plausibly send
- Foundation Edition Appendix D (AI Agent Platforms) connects to advanced civic tech work

---

Assessment Rubric: Civic Issue Brief

Evaluate primarily on whether the brief demonstrates the integrated skills the curriculum has built. Not all sections need to be exemplary; consistency across sections is more important than peak performance in one.

Criterion	Developing	Proficient	Exemplary
Issue is real and specific	Issue is hypothetical or too broad	Real local or state issue, framed as specific question	Real issue with clear stakes for the student's community; question is precisely framed

Criterion	Developing	Proficient	Exemplary
Responsible agency identified	Vague or wrong	Correct agency identified at correct level of government	Agency identified plus relevant officials or specific decision-makers named
Sources are real and varied	Few sources, all of one type, or AI-only	At least 3 sources of multiple types	5+ sources spanning government, news, and at least one other type
Verification done on factual claims	Unverified claims used	Claims cited with sources	Claims cross-checked across multiple sources; uncertainty acknowledged where present
Different perspectives represented	Single perspective only	Multiple perspectives summarized fairly	Perspectives summarized with each side's strongest argument; areas of agreement and disagreement identified
Position is defensible	No position, or position not supported by sources	Position stated; supported by reasoning from sources	Position stated; defended from sources; counterarguments acknowledged
Honest about uncertainty	Pretends certainty about everything	Acknowledges what student doesn't know	Specific articulation of which questions remain open and what would resolve them
AI Use Statement included	Missing or vague	Clear disclosure of AI use	Disclosure includes specific tasks, distinguishes assistance from authorship, confirms substantive work is student's
Citation completeness	Sources mentioned but not fully cited	Sources cited in consistent format	Citations enable a reader to find each source independently

Answer Key / Scoring Notes

Lesson 5 is a project lesson, so there is no single correct issue or position. The answer key is the civic-research structure: real issue, responsible agency, verified facts, multiple perspectives, defensible position, uncertainty, AI disclosure, and citations.

Issue-selection key: - Acceptable issues are real, local or state-connected, researchable, and tied to an identifiable agency, public official, school board, city/county office, or state office. - Too broad: "pollution," "school safety," "traffic," "climate change," "education." - Workable: "Should [county/city] add more crosswalks near [school/street]?", "How is [county] addressing water quality in [local water body]?", "Should the school board revise the phone policy for middle schools?", "What should the city do about flooding near [neighborhood/road]?"

Self-Check answer key: 1. A "real" issue is an actual public problem or decision connected to the student's school, neighborhood, city, county, or state, with real sources and a real responsible decision-maker. 2. A topic is a broad area; a researchable question is specific enough to investigate. Topic: "school traffic." Researchable question: "Should the city add a crossing guard or crosswalk near [school] during morning drop-off?" 3. An agency is a government office or public body responsible for a specific area of public work, such as a school board, city transportation department, county environmental office, Florida Department of Education, or Florida Fish and Wildlife Conservation Commission. 4. Public records matter because they let citizens inspect government information, decisions, agendas, budgets, reports, and correspondence instead

of relying only on summaries or rumors. 5. The AI Use Statement matters because it shows what AI did and did not do, protects academic integrity, and helps readers trust that the student's research, reasoning, and position are their own.

Minimum proficient Civic Issue Brief: - Specific question rather than broad topic. - Correct responsible agency or official. - At least three cited sources, including at least one official or government source where available. - At least three verified facts. - At least two perspectives represented fairly. - A position or honest uncertainty statement supported by sources. - AI Use Statement, even if the statement says no AI was used.

Common grading mistakes to avoid: - Do not grade the student's political or policy position. Grade whether it is defensible from evidence. - Do not accept AI output as a source by itself. AI may help brainstorm, organize, or ask questions; factual claims still need real sources. - Do not require every student to produce a perfect full policy brief in one pilot lesson. Proficient work is specific, verified, and honest, even if compact.

Sample Proficient Artifact (for educator reference)

```
CIVIC ISSUE BRIEF

Student: L.M.
Date: May 15, 2026

Topic:
Traffic safety near school drop-off.

The Question I Researched:
Should the city consider adding a marked crosswalk or crossing guard near the
intersection closest to our school's morning drop-off line?

Why This Issue Matters Locally:
Many students walk or bike near the school entrance while cars are turning into
the drop-off line. The issue matters because it affects student safety and
traffic flow during the same 30-minute window each morning.

Responsible Agency or Official:
The city transportation/public works department appears responsible for local
crosswalks and traffic signs. The school administration and school board may also
be involved because the issue is connected to school arrival procedures.

Key Facts With Sources:
1. The city public works page says residents can request review of traffic safety
   concerns through the city service request form. Source: city public works /
   transportation webpage.
2. The school handbook says morning drop-off begins at 7:45 a.m. and students
   should use designated entrances. Source: school handbook.
3. A local news article about school-zone safety reported that driver speed and
   visibility are common concerns near schools. Source: local news article.

Different Perspectives:
Some parents may support a new crosswalk or crossing guard because they want a
clearer safety procedure. Some drivers may worry that another crossing point will
slow traffic or create confusion. The city may need evidence before changing signs
```

```
or staffing a crossing guard.

What I Think Should Happen and Why:
I think the school should first collect arrival observations for one week and then
send a request to the city transportation department. This is better than asking
for a solution immediately because the city will need specific evidence about
where students cross, how many cars pass, and when the risk is highest.

What I Don't Know:
I do not know whether the city has already studied this intersection. I also do
not know how much a crossing guard costs or who pays for it.

AI Use Statement:
I used AI to brainstorm possible agencies connected to traffic safety and to help
me make my research question more specific. I verified the agency information on
the city website. The writing, position, and source notes are my own.

Sources Cited:
- City public works / transportation webpage, "Traffic Safety Requests."
- School handbook, arrival and drop-off section.
- Local news article on school-zone traffic safety.
```

This sample is **Proficient** because it frames a specific real issue, identifies a plausible agency, uses multiple source types, acknowledges uncertainty, and discloses AI assistance. It becomes **Exemplary** if the student provides exact URLs/titles/dates, uses a school board agenda or city traffic document, includes a stronger counterargument, and identifies the named official or department contact.

Additional Scored Exemplars

Use these examples to calibrate Civic Issue Briefs across levels and accommodations.

Exemplar A — Exemplary

Question: “Should the county add a marked crossing or traffic-calming review near the library entrance used by students after school?”

The student cites a county public works service-request page, a school dismissal policy, a local crash-data dashboard, and a local news article about school-zone traffic. The brief identifies the county public works department and the school principal as separate stakeholders. The position: request a one-week observation study before asking for a permanent traffic change. The brief includes a counterargument: new signage may slow traffic or require funding. The “What I Don’t Know” section names the missing cost estimate and whether the road is city- or county-maintained. AI use is disclosed as brainstorm/search-term support only.

Score: Exemplary. The issue is precise, the agency logic is careful, sources are varied, uncertainty is honest, and the recommendation is proportional to the evidence.

Exemplar B — ELL-adapted Proficient

Question: “Should our city add more trash cans at [local park]?”

The student uses three sources: city parks webpage, a local news article about park cleanup, and an advocacy post from a neighborhood group. The writing uses sentence frames: “The agency is…” “The evidence says…” “I think…” The position is: add two trash cans near the picnic area and post signs before weekends. The student writes: “I do not know the cost. I need city budget information.” AI disclosure: “I used AI to help translate ‘public works’ and ‘parks department.’ I checked the city website.”

Score: Proficient. The language is simple, but the civic reasoning is present: real issue, agency, sources, position, uncertainty, and AI disclosure.

Exemplar C — IEP-modified Proficient

Artifact form: recorded two-minute presentation plus a one-page source checklist.

Question: “Should the school change the bike-rack location?”

The student identifies the school administration and district facilities office, uses a school map, a handbook excerpt, and a student observation tally. The student gives a defensible position orally and submits a checklist showing source titles and what each source proved. The AI Use Statement is read aloud and recorded: “No AI wrote this. I used AI to ask what office handles school facilities, then checked the district page.”

Score: Proficient if oral presentation is an approved accommodation. The rubric criteria are met through an alternate format; do not penalize for not producing a full written brief.

Exemplar D — Beginning / needs reteach

Topic: “climate change is bad.”

The student gives one AI paragraph, no agency, no local/state question, no cited source, and no AI Use Statement.

Score: Beginning. The student has not narrowed the issue, verified claims, identified a responsible agency, or documented AI use. Reteach topic narrowing and require a revised compact brief using a teacher-provided source packet.

Extension Options

Extension 1: Actually send the letter. If a student produced a Letter to a Public Official, encourage them (with parent permission and your guidance) to actually send it. Local officials respond to constituent letters more often than students expect.

Extension 2: Comparative civic project. Have the student do a parallel mini-project on the same issue in a different jurisdiction (a different Florida county, a different state). What’s similar? What’s different? Why?

Extension 3: Public records request. For students particularly interested in their issue, walk them through how to file a Florida public records request for relevant documents. Florida’s Sunshine Law (Chapter 119, Florida Statutes) makes most government records accessible by simple written request — no formal forms required. A request typically identifies the records sought, the requesting party, and a delivery method; agencies must respond within a reasonable time and may charge for staff time and copying.

Extension 4: Foundation Edition reading. Chapters 15 and 16 of *The Open Source Student* (Foundation Edition) cover civic engagement and what students who build can shape. Connects to today's work.

No-Install Adaptation Notes

The mini-project requires source research, which means students need access to information. Two paths:

Path A — Internet access for student devices: The standard implementation. Students search, find sources, work on their briefs.

Path B — Limited or no student internet access: - Provide pre-researched source packets for 5–10 issue options - Each packet contains: a short issue summary, 3–5 cited sources reproduced or excerpted in print, identification of the responsible agency - Students choose from these packets and complete the brief using the provided sources - The work of source-finding is reduced; the work of analysis, position-taking, and brief production is preserved

For pilot adopters in settings without student internet, prepare these source packets in advance. The lesson is fully workable on Path B; only the source-discovery work is reduced.

For AI use during the project: same adaptation as previous lessons — teacher demonstration of an AI service in the role of research assistant; students observe and apply the demonstrated approach to their own briefs without necessarily having direct AI access.

Common Misconceptions

"Civic engagement means political activism." Civic engagement is broader. Researching an issue, writing to an official, attending a public meeting, voting, contributing to community discussion — all of these are civic engagement, regardless of political orientation.

"My issue isn't important enough." Local issues that affect the student's community are exactly the right scale for this project. The point isn't to take on the biggest national issue; it's to do real civic work at a real scale.

"I should be neutral." The brief includes a position because civic literacy includes the capacity to take a defensible position. Neutral journalism is valuable; it is not the only valuable form of civic communication. The brief asks for evidence-based position-taking.

"AI can do the project for me." The lesson and the disclosure requirement make this transparent. AI assistance is fine; AI authorship is not. The student must do the substantive work, including verification, position-taking, and synthesis.

"My teacher will grade me based on what I conclude." No. The grade is on whether the conclusion is defensible from the sources, not on what the conclusion is. Reassure students who are anxious about taking positions on charged topics.

Cross-Curricular Connections

Social Studies (Civics): the lesson is itself civics work in compact form. Connects to virtually every aspect of middle-school civics curriculum.

ELA (Communication — Argument and Research): the brief is substantive argument writing with sources. Directly serves B.E.S.T. ELA Communication-strand standards (note: B.E.S.T. ELA uses "C" prefix for what other frameworks call Writing) — specifically ELA.7.C.1.3 (argument with claim, evidence, counterclaim) and ELA.7.C.4.1 (research drawing on multiple reliable and valid sources).

Social Studies (History): civic issues often have historical context. Coordinate with the social studies teacher for issues with relevant historical background.

Science: environmental and health issues often have substantial scientific content. Coordinate with the science teacher when student topics intersect.

Mathematics: policy issues with quantitative dimensions (budgets, statistics, projections) require students to handle numerical claims. Coordinate with the math teacher for issues with significant quantitative content.

Parent Communication Notes

This lesson can produce strong student work that parents will want to see. It can also produce work on civic issues that may be charged.

A student produces particularly strong work. Share with the parent: "[Student] completed our culminating Civic Tech Mini-Project today on [issue]. The brief demonstrates real research, verification, and analytical thinking. You may enjoy reviewing it together."

A student chooses a topic that may be sensitive in their family. Be neutral. The point is the skill, not the position. If the family asks about a particular topic choice, explain that the curriculum builds source-evaluation and civic-engagement skills regardless of position; specific topic choices are the student's.

A student wants to actually send their Letter to a Public Official. Wonderful. Get parent sign-off. Provide guidance on how to send and what to expect. Real civic engagement at the middle school level is a positive outcome.

Closing the lesson well — and the pilot

This is the last lesson of the Pilot Kit. The closing should mark the achievement of the whole week, not just this lesson.

End-of-curriculum language: "You finished the Florida Pilot Kit. Five lessons. Five skills. Five artifacts in your portfolio. You can explain what AI is, you can verify, you can read across sources, you can use AI ethically, and you can do real civic work with all of those skills together. That is a real foundation. You will use it across your education and across your life."

If feasible, take a moment to acknowledge the students individually — even briefly. The Pilot Kit is intentionally short; the relationships students built with the material this week are part of what will determine whether they keep using these skills.

Then transition to the pilot feedback. Direct students to know that their feedback (and their educator's feedback) shapes the next version of the curriculum. They are not just consumers of this work; they are contributors to what comes next.
