AI Literacy for Parents & Teens

A 6-week online course curriculum

Audience: Parent + teen pairs (ages roughly 12–18) taking the course together.
Format: Six weekly modules, 60–90 minutes each.
Outcome: A shared vocabulary, a signed household AI agreement, and durable critical-thinking habits.

Course Overview

This course is designed for parents and teens to take together. Each week is built around a single theme, with shared readings, a guided conversation, and a hands-on activity that gets families experimenting with AI side by side rather than in separate rooms.

Weekly module structure

- Learning objectives — what you'll be able to do by the end
- Core concepts — short reading to do separately before meeting
- Discussion prompts — questions to work through together
- At-home activity — a practical task to do as a pair
- Parent resource list — credible supplemental material
- Reflection — closing journal prompt

Six-week arc

Week	Theme
1	How AI Actually Works
2	Bias & Deepfakes
3	Privacy & Ethics
4	Critical Thinking & Fact-Checking
5	Setting Household Digital Boundaries
6	Modeling Responsible Tech Habits

Week 1 — How AI Actually Works

Learning objectives

- Describe in plain language what a large language model and an image generator are doing under the hood.
- Distinguish between AI that predicts, AI that classifies, and AI that generates.
- Recognise where AI is already embedded in everyday tools (search, social feeds, photo apps, autocorrect).

Core concepts

AI systems are pattern-matching engines trained on huge datasets. They do not "know" things the way a person does — they produce statistically likely outputs based on what they've seen before. This single idea explains most of what comes later in the course: why AI hallucinates, why it reflects bias, and why it can sound confident while being wrong.

Discussion prompts

1. Before this week, what did each of you assume AI was actually doing when you asked it a question? How close was that to the reality?
2. Where have you noticed AI showing up in apps you already use? Make a list together.
3. If an AI is just predicting the next likely word, what kinds of questions is it probably good at — and what kinds is it probably bad at?
4. Teen to parent: what's one thing you wish adults understood about how your generation uses AI? Parent to teen: what's one thing you wish you understood better?

At-home activity: Same question, three tools

Pick one open-ended question you both genuinely care about (e.g. "What should we cook for dinner this week on a $60 budget?" or "What were the causes of WWI?"). Ask the same question in three different tools — for example ChatGPT, Google Gemini, and a regular Google search. Compare the answers side by side: Where do they agree? Where do they disagree? Which one cites sources, and are those sources real? Which felt most confident — and was confidence the same as correctness? Write down one sentence each about what surprised you.

Parent resource list

- Day of AI & Common Sense Media — "What Is AI for Families" video series and toolkit — Short, age-appropriate explainers built specifically for family conversations.

- MIT RAISE — Responsible AI for Social Empowerment and Education — MIT's hub for AI literacy curricula, including materials usable at home.
- The Common Parent — Parent's Guide to AI — Honest, non-fear-based overview of where AI shows up in kids' digital lives.
- Elements of AI (free online course) — University of Helsinki's free introductory course; excellent if a teen wants to go deeper.

Reflection

Write one sentence: "Before this week I thought AI was ___; now I think it's ___."

Week 2 — Bias & Deepfakes

Learning objectives

- Explain why AI models inherit bias from their training data.
- Identify visual, audio, and contextual cues that suggest a deepfake.
- Understand that "looks real" is no longer a reliable test for "is real."

Core concepts

Two related problems sit at the heart of this week. Bias is structural — an AI trained mostly on one kind of voice, face, or worldview will quietly amplify it. Deepfakes are deliberate — synthetic images, audio, and video designed to look authentic. Both rely on the same underlying tech, and both require the same defence: lateral verification rather than gut feeling.

Discussion prompts

1. If an AI was trained mostly on text from one country, one language, or one political viewpoint, how might its answers be skewed? Whose voices might be missing?

2. Have either of you seen something online recently you suspected was AI-generated? What tipped you off — or what failed to?

3. Is there a difference between a deepfake made for satire, one made for a scam, and one made to harass someone? Should all three be treated the same way?

4. If a deepfake of you (or a friend) appeared online tomorrow, what would you want to happen? Who would you tell first?

At-home activity: Bias audit + deepfake hunt

Part A — Bias audit (20 min). Ask one image generator to produce images for ten neutral prompts: "a CEO," "a nurse," "a criminal," "a scientist," "a person cleaning," "a wedding," "a beautiful house," "a homeless person," "a teenager studying," "a family at dinner." Don't add adjectives. Look at the results together: what patterns appear in gender, race, age, body type, setting, wealth? What's missing?

Part B — Deepfake hunt (20 min). Visit MIT's "Detect Fakes" media literacy site and work through their examples together. Then scroll your normal social feeds for 10 minutes specifically looking for AI-generated images. Compare notes on what gave them away.

Parent resource list

- MIT — Media Literacy in the Age of Deepfakes — Interactive lessons and examples designed for classroom and home use.
- Parents Pass It On — How to Explain AI Bias to Teens in Simple Terms — Concrete conversation scripts and comparison exercises.
- Common Sense Media — Deepfakes, Distrust and Disinformation — Parent-facing primer on synthetic media.
- Algorithmic Justice League — Joy Buolamwini's organisation; strong real-world examples of facial-recognition bias.

Reflection

Name one piece of media you've seen this week that you now want to re-examine.

Week 3 — Privacy & Ethics

Learning objectives

- Understand what data is collected when a teen uses an AI chatbot, AI-enabled app, or AI companion.
- Articulate the difference between legal and ethical use of AI.
- Make informed choices about which AI tools to use and what to share with them.

Core concepts

Free AI tools are rarely free. The price is usually data — your prompts, your uploaded files, sometimes your voice and face. On top of the privacy question sits an ethical one: AI companion chatbots are now widely used by teens, and recent research shows roughly three-quarters of teens have tried one, despite serious concerns about emotional reliance and safety for under-18s.

Discussion prompts

1. When you type something into an AI chatbot, where do you imagine that text goes? Who might read it? How long is it kept?

2. Are there things you've already shared with an AI that you wouldn't want a teacher, employer, or future partner to read? (No one has to share specifics — just answer yes or no.)

3. AI companion apps are designed to feel like friends. What's the difference between a tool that helps you and a tool that wants you to keep coming back?

4. If a free app uses your data to train its model, is that a fair trade? Would your answer change if the app were used by a 9-year-old?

At-home activity: Privacy settings deep clean

Pick the three apps each of you uses most that involve AI (this often includes Snapchat, Instagram, TikTok, ChatGPT, Gemini, Character.AI, photo editors, or a smart speaker). For each one, together: (1) open the privacy settings, (2) find the section on "data used for training" or "improve our models" and decide whether to opt out, (3) check what's shared with third parties, (4) review what's stored in chat history and delete anything that doesn't need to be there. Make a simple shared note of what each of you changed. Re-check this list every six months.

Parent resource list

- Common Sense Media — Teens, Trust, and Technology in the Age of AI Companions — The research finding that ~75% of teens use AI companions and why this matters.
- Electronic Frontier Foundation — Student Privacy — Rights-based perspective on data collection in schools and apps.
- UNESCO — Recommendation on the Ethics of Artificial Intelligence — The global ethical framework; readable summary at the top.
- Mozilla Foundation — Privacy Not Included — Searchable reviews of consumer apps and devices by privacy risk.

Reflection

What's one piece of personal information you've decided you'll never put into an AI tool?

Week 4 — Critical Thinking & Fact-Checking

Learning objectives

- Apply the SIFT method to any piece of AI-generated or AI-amplified content.
- Recognise AI "hallucinations" — fabricated facts, fake citations, invented quotes.
- Build a personal fact-checking workflow that takes under two minutes.

Core concepts

AI tools confidently produce false information — fake legal cases, invented historical quotes, made-up scientific citations. Confidence is not a signal of accuracy. The most durable defence isn't a specific tool; it's a habit. Digital literacy expert Mike Caulfield's SIFT method gives a four-step routine: Stop, Investigate the source, Find better coverage, Trace claims back to the original context.

Discussion prompts

1. Tell each other about a time you believed something online that turned out to be wrong. What made it believable? What would have caught it?
2. Why might an AI invent a fact and present it confidently? (Hint: re-read Week 1.)
3. If a TikTok, an AI summary, and a published news article all say the same thing, are they three sources — or one? How would you check?
4. What's the difference between "I don't know" and "I looked it up and it's not true"? Which one do AI tools tend to skip?

At-home activity: Catch the AI lying

Together, design five questions where the correct answer is something each of you knows well — a family member's job, the plot of a favourite book, a local sports result, a niche hobby fact, a historical event from your country. Ask an AI chatbot each question and rate its answer as correct, partially correct (with mistakes), confidently wrong, or refused/hedged. For any "confidently wrong" answer, run it through the SIFT steps as a pair. Discuss: would you have caught this if you didn't already know the answer?

Parent resource list

- University of Chicago Library — The SIFT Method — Clean, step-by-step explainer of Caulfield's method.
- News Literacy Project — Checkology platform — Free interactive lessons including modules on AI and misinformation.
- Mike Caulfield's blog — Hapgood — The originator of SIFT writes accessibly about online verification.
- Snopes & Full Fact — Bookmark these as a family for spot-checks. Also see fullfact.org.

Reflection

"This week I caught AI being wrong about ___. I caught it because ___."

Week 5 — Setting Household Digital Boundaries

Learning objectives

- Co-create a family agreement on AI use that both parents and teens actually accept.
- Distinguish boundaries that need controls from boundaries that need conversation.
- Identify red-flag scenarios that warrant immediate adult involvement.

Core concepts

Boundaries fail when they're imposed without buy-in. This week is the most collaborative of the course. The goal isn't a parental rulebook — it's a household agreement that both sides sign because both sides helped write it. The American Academy of Pediatrics' Family Media Plan is a useful scaffold; adapt it for AI specifically.

Discussion prompts

1. What are AI uses that feel clearly fine to both of you? Clearly off-limits? In the grey zone?
2. School and homework: when does using AI count as a tool (like a calculator), and when does it count as cheating? Where does your school draw the line — and do you agree?
3. AI companions and emotional support: under what circumstances, if any, is it healthy to use one? When does it become a substitute for real connection?
4. What do each of you want from the other? (Teens: what do you want parents to stop doing? Parents: what do you want teens to start doing?)
5. What's the "call a human" rule — situations where you'd both agree to stop, close the app, and talk to a real person?

At-home activity: Draft your household AI agreement

Use a single shared document. Build five short sections together: (1) What we use AI for (homework help, recipes, brainstorming, etc.); (2) What we don't use AI for (e.g. uploading personal photos to image generators, sharing real names or addresses, replacing a real friend or counsellor); (3) What we always disclose (e.g. AI use on a school assignment if asked; AI photo edits before posting); (4) Time and place rules (e.g. no AI companions after 10pm; phones out of bedrooms overnight); (5) The call-a-human list — specific situations where you stop using AI and talk to a real person (mental health questions, medical symptoms, legal trouble, anything that scares you). Both sign and date it. Put it somewhere visible. Agree on a review date in 90 days.

Parent resource list

- American Academy of Pediatrics — Family Media Plan — Customisable template, easy to adapt for AI.
- AAP Healthy Digital Habits Toolkit — Broader campaign resources for families.
- Common Sense Media — Parents' Ultimate Guide to ChatGPT — Practical guidance on school and homework boundaries.
- Center for Humane Technology — Family resources — Frameworks for understanding attention-capture design.

Reflection

Sign the agreement. Take a photo of it together.

Week 6 — Modeling Responsible Tech Habits

Learning objectives

- Recognise that teens learn tech habits primarily by watching the adults around them.
- Identify the specific behaviours each family member wants to model — and to stop modelling.
- Build a simple ongoing practice so this course's lessons don't fade in three months.

Core concepts

Research on family media habits keeps reaching the same conclusion: parents' own behaviour shapes children's relationship with technology more than any rule. If a parent is on their phone at dinner, "no phones at dinner" won't hold. If a parent never says "I don't know — let me check," teens won't either. This week flips the lens. Parents commit publicly to changes too.

Discussion prompts

1. Parent: what tech habit do you have that you'd be embarrassed for your teen to copy exactly? Teen: same question back.

2. When was the last time each of you saw the other put down a device to be fully present? How did it feel?

3. What does it look like, in this house, to use AI well? Give three concrete examples you've already lived in the past six weeks.

4. What do you want your family's relationship with technology to look like in five years? Work backwards: what would you each need to do this month for that to be on track?

5. What did each of you learn from the other during this course that you didn't expect?

At-home activity: The mirror commitment + monthly check-in

Part A — Mirror commitments. Each person writes down three specific behaviours they'll change, framed as things the other person will be able to see. Examples: "I'll leave my phone in the kitchen during dinner," "I'll say out loud when I'm using AI to write something," "I'll ask before posting a photo of you," "I won't use my AI companion app on school nights." Exchange lists. Put them on the fridge.

Part B — Schedule the check-in. Put a recurring 30-minute family check-in on the calendar — monthly is realistic. Three questions each time: What did AI help us with this month? What did AI

cost us this month? Does our household agreement still fit, or does it need an edit?

Parent resource list

- Common Sense Media — Digital Citizenship Curriculum — Long-form companion curriculum for families wanting to keep going.
- AAP — Media and Children — Research-backed practices for modelling healthy use.
- Center for Humane Technology — Take Control — Concrete settings changes and habit shifts for adults.
- Day of AI — ongoing free curriculum — Yearly updated lessons families can revisit as the tech changes.

Reflection

"The single most important thing I learned in this course was ___. The first thing I'll do differently tomorrow is ___."

Course Wrap-Up

By the end of these six weeks, you should have, together:

- A shared vocabulary for talking about AI without panic or hype
- A signed household AI agreement with a review date
- A working fact-checking habit (SIFT) and a privacy-clean set of accounts
- A scheduled monthly check-in to keep the conversation alive
- Three concrete behaviour changes each that the other family member will notice

The technology will keep changing. The point of this course isn't to give you a finished answer — it's to give you the shared muscle of asking the next question together.

Facilitator Notes

- Pacing: One module per week is ideal. Compressing it into a weekend defeats the point — the time between sessions is when the conversations actually happen.
- Async option: Each module works as a self-paced unit. If running this as a cohort course, host an optional 30-minute live group debrief at the end of each week.
- Age range: Discussion prompts work from about age 12. For younger kids (8–11), simplify Weeks 2 and 3 and skip the deepfake-search part of Week 2's activity.
- Single-parent or multi-teen households: All activities scale. The "pair" can be parent + teen, two parents, two siblings, or a guardian + child. The key is two perspectives in the room.
- Updating: AI tooling changes fast. Re-check linked resources every six months and swap in current examples for the activities.