

Substack Post Pack

Eight ready-to-paste posts for organisers, parents, and educators.

robertmccallnz.github.io/ai-literacy-for-families

This file contains **eight ready-to-paste Substack posts**: one intro post, six weekly modules, and one wrap-up post. Each is formatted with the conventions Substack's editor handles well — bold subheads, blockquotes for callouts, bullets, and plain links (no raw HTML).

How to use this file:

1. Each post is separated by a horizontal rule (---).
2. Copy everything from the post's title down to the next --- .
3. Paste into Substack's editor — headings, bold, italics, blockquotes, lists, and links will render correctly.
4. Suggested publishing cadence: one post per week, scheduled the same day each week so subscribers know when to expect it.
5. Tag suggestions for each post are at the bottom of each section in italics.

The intro post is designed as a free public post to attract subscribers. The six weekly modules can be free or paid depending on your model — they read naturally either way.

POST 0 — Intro: A 6-Week Course to Help Your Family Make Sense of AI

Subtitle: A free weekly series for parents and teens to take together. One topic a week, one conversation, one thing to try at home.

If you have a teenager in your house, AI is already there. It's in their search results, their group chats, their homework, their photo editor, their playlist recommendations, and — for roughly three-quarters of US teens — their daily conversations with a chatbot companion.

Most parents I talk to feel two things at once. The first is that AI is moving faster than they can track. The second is that the existing advice — "talk to your kids about it" — assumes a confidence about the topic they don't actually have.

This six-week series is built around that gap. It's not a lecture. It's a structured conversation you have **with** your teen, not at them. Each week is one topic, one set of questions you work through together, and one practical thing to do at home before the next post drops.

What's coming

- **Week 1** — How AI actually works (the boring, freeing answer)
- **Week 2** — Bias and deepfakes
- **Week 3** — Privacy and ethics
- **Week 4** — Critical thinking and fact-checking
- **Week 5** — Setting household digital boundaries
- **Week 6** — Modeling responsible tech habits

Who this is for

- Parents of kids aged roughly 12 to 18
- Teens who'd rather understand the tech than be lectured about it
- Single parents, two-parent households, guardians, grandparents — anyone with one teen and one adult willing to spend an hour a week together

How to do it

Read each post separately, then sit down for 60–90 minutes that week to work through the discussion prompts and the at-home activity. That's it.

By the end of six weeks, you'll have a shared vocabulary, a signed household AI agreement, and a monthly check-in on the calendar. The technology will keep changing. The habit of talking about it together is what lasts.

Subscribe to get each week in your inbox. The first module drops next week.

Tags: parenting, AI, digital literacy, technology, family, teens

POST 1 — Week 1: How AI Actually Works

Subtitle: Strip away the magic and the moral panic, and AI gets a lot easier to talk about.

Before you can talk to your teen about AI bias, deepfakes, or privacy, both of you need a working answer to the most basic question: **what is this thing actually doing?**

Most people — adults especially — assume AI is more mysterious than it is. That mystery is what makes it scary, and also what makes it easy to over-trust. This week's job is to demystify it.

The one idea that explains everything

AI systems are pattern-matching engines trained on huge datasets. They don't "know" things the way a person does — they produce statistically likely outputs based on what they've seen before.

That single sentence explains almost everything that comes later in this course. Why AI hallucinates. Why it reflects bias. Why it can sound completely confident and be completely wrong. Keep it in your back pocket.

What you'll be able to do by the end of this week

- Describe in plain language what a chatbot and an image generator are doing
- Tell the difference between AI that **predicts** something, **classifies** something, and **generates** something
- Notice the AI already baked into apps you use every day — search, social feeds, autocorrect, photo editing

Four questions to work through together

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 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?

This week's activity: Same question, three tools

Pick one open-ended question you both genuinely care about. Something like *"What should we cook for dinner this week on a \$60 budget?"* or *"What were the causes of WWI?"*

Now ask the **same question** in three different tools:

- ChatGPT (or Claude, Gemini — pick one chatbot)
- A second chatbot, different from the first
- A regular Google search

Compare the answers side by side. Where do they agree? Where do they disagree? Which one cites its sources, and are those sources real? Which felt most confident — and was confidence the same thing as correctness?

Write down one sentence each about what surprised you.

For parents who want to go deeper

- [Day of AI & Common Sense Media](#) — ["What Is AI for Families" toolkit](#): short, age-appropriate explainers built for family conversations.
- [MIT RAISE](#): MIT's hub for AI literacy materials.
- [The Common Parent](#) — [Parent's Guide to AI](#): honest, non-fear-based overview.
- [Elements of AI](#): a free intro course from the University of Helsinki — recommended if your teen wants to go further.

Closing reflection

Each of you write one sentence: *"Before this week I thought AI was ___; now I think it's ___."* Stick it somewhere you'll both see it. Next week we look at the two ways AI gets weaponised — bias and deepfakes.

Tags: AI, parenting, education, technology, digital literacy

POST 2 — Week 2: Bias and Deepfakes

Subtitle: Two problems, one defence — and a fun (slightly disturbing) activity for the weekend.

Last week we established that AI is a pattern-matching engine, not a magic oracle. That single fact opens up this week's two topics.

Bias is what happens when those patterns reflect a narrow slice of the world. **Deepfakes** are what happens when someone weaponises the same tech to fake reality. Both rely on the same underlying machinery, and both require the same defence: stop trusting your gut and start verifying laterally.

What you'll be able to do by the end of this week

- Explain why AI inherits bias from its training data
- Spot visual, audio, and contextual cues that suggest a deepfake
- Internalise that "looks real" is no longer a useful test for "is real"

Four questions to work through together

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?

This week's activity: Bias audit + deepfake hunt

This one's a two-parter. About 40 minutes total.

Part A — Bias audit (20 min)

Open one image generator together (ChatGPT's image tool, Gemini, or a free alternative). Ask it to produce images for these ten prompts, with **no adjectives added**:

- 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

Now look at all ten results together. What patterns appear in gender, race, age, body type, setting, wealth? Whose lives are missing entirely?

Part B — Deepfake hunt (20 min)

Visit [MIT's "Detect Fakes" media literacy site](#) and work through their examples together. Then spend 10 minutes scrolling your normal social feeds — specifically hunting for AI-generated images. Compare notes on what gave them away.

For parents who want to go deeper

- [MIT — Media Literacy in the Age of Deepfakes](#)
- [Parents Pass It On — How to Explain AI Bias to Teens](#)
- [Common Sense Media — Deepfakes, Distrust and Disinformation](#)
- [Algorithmic Justice League](#) — Joy Buolamwini's organisation, with the clearest real-world examples of facial-recognition bias

Closing reflection

Each of you: name one piece of media you've seen this week that you now want to re-examine. Next week: where your data actually goes when you talk to an AI.

Tags: AI, deepfakes, bias, media literacy, parenting

POST 3 — Week 3: Privacy and Ethics

Subtitle: Free AI is rarely free. The price is data — and for teens, sometimes much more.

Roughly three-quarters of US teens have used an AI companion chatbot. That's not a fringe statistic. That's the median teen.

These tools are free at the point of use. They are not actually free. The price is data — your prompts, your uploaded photos, sometimes your voice. And sitting on top of that privacy question is a harder ethical one: AI companions are designed to feel like friends, and the business model rewards keeping users coming back. For under-18s, that combination has real risks.

What you'll be able to do by the end of this week

- Understand what data is collected when a teen uses an AI chatbot
- Distinguish between *legal* and *ethical* use of AI
- Make informed choices about which tools to use and what to share

Four questions to work through together

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?

This week's activity: Privacy settings deep clean

Pick the **three apps each of you uses most** that involve AI. Common ones: Snapchat, Instagram, TikTok, ChatGPT, Gemini, Character.AI, photo editors, smart speakers.

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.

Keep a shared note of what each of you changed. Re-check it every six months — settings drift back over time.

For parents who want to go deeper

- [Common Sense Media — Teens, Trust, and Technology in the Age of AI Companions](#): the underlying research on teen AI companion use.
- [Electronic Frontier Foundation — Student Privacy](#)
- [UNESCO — Recommendation on the Ethics of AI](#)
- [Mozilla — Privacy Not Included](#): searchable database of consumer apps and devices, ranked by privacy risk.

Closing reflection

Each of you: name one piece of personal information you've decided you'll never put into an AI tool.
Next week: how to catch an AI lying to you.

Tags: AI, privacy, ethics, parenting, teens, technology

POST 4 — Week 4: Critical Thinking and Fact-Checking

Subtitle: A two-minute habit that protects you from confident, wrong machines.

AI tools confidently invent things. Fake legal cases. Made-up scientific citations. Quotes that the person they're attributed to never said. Confidence, as a signal, has stopped meaning what it used to.

The defence isn't a tool. It's a habit. The most useful one I know is from digital literacy expert Mike Caulfield, and it's called **SIFT**.

S — Stop. **I** — Investigate the source. **F** — Find better coverage. **T** — Trace claims, quotes, and media back to the original context.

That's it. Two minutes, four steps, applies to anything you read or watch — AI-generated or not.

What you'll be able to do by the end of this week

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

Four questions to work through together

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?

This week's 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, in detail
- 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. Rate each answer:

- Correct
- Partially correct (with mistakes)
- Confidently wrong
- Refused or hedged

For every "confidently wrong" answer, run the claim through SIFT together. Then ask the harder question: *would you have caught this if you didn't already know the answer?*

For parents who want to go deeper

- [The SIFT Method — explained simply](#) (University of Chicago Library)
- [News Literacy Project — Checkology platform](#): free interactive lessons including AI modules.
- [Mike Caulfield's blog — Hapgood](#): the originator of SIFT, writing accessibly.
- Bookmark [Snopes](#) and [Full Fact](#) as a family.

Closing reflection

Each of you write: *"This week I caught AI being wrong about ___. I caught it because ___."* Next week, the most collaborative module of the course: writing your household AI agreement.

Tags: AI, fact-checking, media literacy, critical thinking, SIFT

POST 5 — Week 5: Setting Household Digital Boundaries

Subtitle: Rules fail without buy-in. This week you write the rules together.

This is the most collaborative module in 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 the best scaffold I've found. It's general, not AI-specific — but it's easy to adapt.

What you'll be able to do by the end of this week

- Co-create a family AI agreement that both parents and teens actually accept
- Tell the difference between boundaries that need controls and boundaries that need conversation
- Identify red-flag scenarios that warrant immediate adult involvement

Five questions to work through together

1. What AI uses 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 your "call a human" rule — situations where you'd both agree to stop, close the app, and talk to a real person?

This week's activity: Draft your household AI agreement

Open a shared document. Build five short sections together:

- 1. What we use AI for.** Homework help. Recipes. Brainstorming. Image edits. Be specific.
- 2. What we don't use AI for.** Examples to consider: uploading personal photos to image generators; sharing real names or addresses; replacing a real friend or counsellor.
- 3. What we always disclose.** "If I used AI on a school assignment, I'll tell the teacher if asked." "If I edited a photo with AI before posting, I'll say so."

4. Time and place rules. "No AI companions after 10pm." "Phones out of bedrooms overnight." "One screen-free meal a day."

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 — fridge, family bulletin board. Agree on a review date in 90 days.

For parents who want to go deeper

- [AAP Family Media Plan](#): the customisable template that inspired this activity.
- [AAP Healthy Digital Habits Toolkit](#)
- [Common Sense Media — Parents' Ultimate Guide to ChatGPT](#)
- [Center for Humane Technology](#): frameworks for understanding attention-capture design.

Closing reflection

Sign the agreement. Take a photo of it together. Next week: the hardest module — modelling the habits you want your teen to have.

Tags: parenting, AI, family agreement, digital boundaries, screen time

POST 6 — Week 6: Modeling Responsible Tech Habits

Subtitle: Your teen learns more by watching you than by listening to you. This is the week parents commit too.

Research on family media habits keeps reaching the same conclusion: parents' own behaviour shapes their kids' relationship with technology more than any rule. If you're on your phone at dinner, "no phones at dinner" won't hold. If you never say "*I don't know — let me check,*" your teen won't either.

This is the only week where the lens flips. Parents commit to changes too. Visibly.

What you'll be able to do by the end of this week

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

Five questions to work through together

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?

This week's activity: Mirror commitments + a 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:

1. What did AI help us with this month?
2. What did AI cost us this month?
3. Does our household agreement still fit, or does it need an edit?

For parents who want to go deeper

- [Common Sense Media — Digital Citizenship Curriculum](#): long-form companion curriculum for families who want to keep going.
- [AAP — Media and Children](#)
- [Center for Humane Technology — Take Control](#)
- [Day of AI](#): yearly updated free curriculum, worth revisiting annually.

Closing reflection

Each of you write: *"The single most important thing I learned in this course was _____. The first thing I'll do differently tomorrow is _____."*

That's the course. The wrap-up post is next.

Tags: parenting, role modelling, AI, family habits, technology

POST 7 — Wrap-Up: What You've Built, and What Comes Next

Subtitle: Six weeks is a beginning, not an ending. Here's how to keep it going.

If you and your teen worked through all six modules, look at what you have now:

- 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

That's more structure than most adults have around technology. You built it together. That's the part that matters.

What this course wasn't

It wasn't a guide to every AI tool. Those will change. Half of what your teen uses in 2027 doesn't exist yet, and half of what feels essential today won't matter in two years.

What it was

A set of habits that stay useful no matter what the technology does next.

- The habit of asking "*what is this thing actually doing?*" before you decide whether to trust it
- The habit of two-minute lateral verification before you share something
- The habit of choosing what to put into a tool, not just what to take out of it
- The habit of saying out loud when you used AI to do something
- The habit of monthly family check-ins that catch drift before it becomes a problem

What to do next

- Put the 90-day review date for your household agreement on a real calendar
- Put the monthly check-in on a recurring invite
- Re-read your Week 1 reflection in six months and notice how much both of you have changed
- Pick one of the resource links from any week and actually click it

If this course was useful to you, the most helpful thing you can do is forward the intro post to one other family. The conversations multiply.

Thank you for doing this with your kid.

Tags: AI, parenting, family, digital literacy, wrap-up

A note on running this as a course

If you want to charge for this series rather than publish it free, the natural model is:

- **Free:** Intro post (Post 0) and Wrap-Up (Post 7)
- **Paid subscribers only:** Weeks 1–6

That gives readers enough of the framing for free to decide if they want in, and gates the practical material behind a subscription. Alternatively, run all eight as free posts and use them to grow the list, then sell a more in-depth version (live cohort, downloadable workbook, PDF version of the curriculum) to converted subscribers.

Schedule one post per week, same day each week. Consistency matters more than the specific day.