

The Conscious Parent's Guide to AI

BY JORDAN URBS

81% of parents have already used AI for something parenting-related.

And 52% of Americans say asking AI for parenting advice is a bad idea.

Which means most of us are doing this quietly, and nobody's comparing notes.

I'd rather compare notes.

I'm a dad who builds with AI for a living. My kids and I make bedtime books, videos, and video games together with it. They come up with the characters and plots, and I help manage the tech part. They point out the mistakes to be fixed.

I've spent 3 years now figuring out which parts of this technology belong in my house, and which parts don't.

This guide is the notes. Three parts:

1. **Keep it private** — how to ask AI your real parenting questions without handing your family's inner life to a big tech company
2. **Keep your values** — how to use AI without outsourcing the one job that's yours (parenting)
3. **Build something together this week** — three projects you can do with your kid today, with the exact prompts

No hype, no doom.

Fellow parents, comparing notes.

Keep it Private

Think about the questions you'd only ask an AI at 11pm, as a parent.

"My 8-year-old is still wetting the bed and I don't want to make him feel bad."

"How do I explain a divorce to a 5-year-old?"

"My daughter got a diagnosis today and I don't understand the report."

Take into account where those questions go when you type them into a standard chatbot account: onto a company's servers, attached to your name and email, preserved forever, sold to advertisers when deemed appropriate, and often used to train the company's next model.

Now, these questions may not seem like a big deal.

But when these types of questions are asked on a consistent basis, and linked to your name & account, it begins to paint a picture of you and your family: your concerns, your lifestyle, your psychology.

These companies use this data to build profiles of you and your family.


And these profiles are not simply sold for advertising, but handed over to **mass surveillance** firms, **law enforcement** and (albeit less intentionally) to hackers and the dark web when the next big data leak occurs (**happens** all the time).

Now imagine someone with the wrong intentions gets their hands on personal details about your family.

I won't go any farther because I'm trying to offer solutions here, not scare you (though I hope by now you understand why this is important).

Fortunately, there are easy solutions to this threat vector.

The fix has three levels for any level of tech competency:



You don't need all three (most families will be fine at level two):

LEVEL 1 **Awareness** 10 minutes, free

Keep using whatever you use (ChatGPT, Claude, Gemini), but deliberately.

- Sort your questions into two piles:
 - logistics (meal plans, packing lists, homework help)
 - *family-sensitive* (health, behavior, emotions, anything with your kid's name/personality in it)
- Logistics can go anywhere. Sensitive questions wait for level 2 or 3.
- In your account settings, turn off "improve the model with my data" (or the equivalent). Most major providers bury this toggle somewhere, and some don't offer it at all. And it doesn't mean they'll actually do it... but it's better than nothing.
- Never put your child's full name, school, or photos into any chatbot, whether you're logged in or not. Always remain ambiguous with your questions concerning your family.

That's level 1.

You've already cut the data trail by most of what matters.

RELATED TIP: Did you know Meta (Facebook, Instagram, etc) by default trains its AIs on your posted photos? You can shut this off deep in the settings but as a rule of thumb it's probably a good idea to simply not post content of your family at all.

LEVEL 2**Private Access**

\$18/month, 10–20 minutes of setup

There are services that sit between you and the big AI models, so the model provider sees the *service* instead of *you*.

Venice is the one I use (\$18/month—not a referral link): the same top-tier models, without your name attached to your questions. There is also **Maple**, **PPQ**, and plenty more popping up lately (because it's becoming such an important topic).

Here's an analogy: imagine you're paying at a store with cash instead of by card. Same purchase, no receipt tied to your name.

This is where the sensitive pile goes. The bed-wetting question, the diagnosis report, the "am I handling this right" question at 11pm. You get real answers from a strong model without building a permanent record of your family's hardest moments on a server tied to your identity.

One honest caveat: generally speaking, you're still trusting the company in the middle to do what it says (less exposure, not zero).

However, with the services mentioned above (and others), there are options to communicate with AI using **E2EE** and **TEE**. Without getting too deep into tech-speak, this basically means that your questions to AI are locked in a box (E2EE) during the trip to the server, and then only unlocked inside a magical, invisible room (TEE) to get an answer. This ensures that neither the internet company nor the AI server operator nor any one who tries to intercept your chats can ever see your private data, making it the most secure way to use AI for sensitive data.

That's level 2. Passable and good enough for most; but if you're ready to geek out for an afternoon, Level 3 awaits:

LEVEL 3**Your Own Machine**

a weekend of setup, hardware you may already own

Free, open weight AI models (e.g. Kimi, Qwen, Gemma) can run on a reasonably modern computer. A recent Macbook Pro or a gaming PC can most likely handle the small-to-mid size ones fine.

Nothing leaves the house.

No account, no logs, no company in the middle.

For a family, this is the deep end of ownership.

The honest trade: these models are noticeably less capable than the frontier models (the newest, biggest ones). Especially so if you're running them on consumer hardware like your home computer.

For bedtime stories, homework explanations, meal plans, and most everyday questions... you won't care.


For genuinely hard questions (or for using AI to work for you)...you will.

And it costs a real weekend.

Tools like **Ollama** and **LM Studio** have made it drag-and-drop easy compared to 2 years ago, but "easy" still means that weekend of downloading models and finding out which ones your machine can run.

The idea underneath all three levels

There's a name for what you just climbed: the trust spectrum.



At one end, you're relying on a company's promise: terms of service, a privacy policy, a settings toggle. The company can rewrite the promise tomorrow, and you're not the one holding the pen.

At the other end, you're relying on things you own: files on your machine, a model running on your own computer (no internet required). There's no promise to break.

Level 1 manages the promise.

Level 2 shrinks what it covers.

Level 3 swaps it for hardware in your house, where there's nothing left to promise.

The reason to learn this as a parent: the spectrum applies to far more than chatbots.

"Who am I trusting with this, and what data is stored about my family if they change their policy?"

Think about the homework app, the school's AI tutor, the baby monitor with its cloud account, the free game asking for microphone permissions.

The honest answer is usually "I have no idea what would happen," but noticing that is most of the skill.

You don't have to own everything... you have to know where you're standing, and to trust on *purpose* instead of *by default*.

The pattern

Use each level for what it's for.

Everyday stuff wherever it's convenient.

Sensitive stuff through private access or your own machine.

That routing habit, more than any single tool, is what keeps you and your family's inner life yours.

Is my own setup exactly right? I don't know. I'm still trusting third parties a lot more than I'd like.

But we're doing our best. *"Where do our questions go?"* is now a question we ask on purpose, and that alone changes what we type.

Keep Your Values

The privacy problem has settings and tools.

But the problem of values doesn't.

In this section we'll discuss a mindset shift required for parents & families in the age of AI. I'll share it from the perspective of a single value.

After 18 months of building non-stop with AI, when it comes to harnessing machines to enhance the lives & upbringing of my children, I keep coming back to this framework:

AI generates options. It cannot decide what's best.

It will produce ten discipline strategies, five ways to explain death to a child, three bedtime stories. All the homework.

And even if it has a deep private "knowledge base" of your child's psychology, it still has no idea which one is right for *your* kid, because "right" isn't a computation.

It's a judgment call.

The deciding is your job. It was always your job. AI makes it possible to hand the job away without noticing—and this is what's dangerous.

So one value-oriented rule in our house: **AI drafts, parents decide.**
Examples:

The machine can propose the chore chart. It doesn't rule on whether the chore chart is fair... or final.

It can summarize the school's reading list; what our family thinks about the list is our business.

And it can generate the apology letter after a sibling fight, but the kid has to rewrite it in his own words, because the point of an apology is that it costs you something.

Why this matters more than screen time

It's the screen-time question: *How much AI should my kid get access to?*

I think that's actually the smaller question.

The bigger one: *How much dependency is AI creating in my child?*

Let's take homework.

A machine can now copy-paste most of what school calls "achievement." The five-paragraph essay, the book report, the practice worksheet: a model produces all of it in seconds, virtually free of cost.

Whatever a machine can copy is getting cheap fast.

But it can't copy judgment, taste, discipline, following through, or simply being the kind of person others trust.

AI has revealed the true purpose behind mundane homework assignments. As parents we must now shift our focus to the traits & skills that school tasks once helped develop in children, instead of the tasks themselves.

So when I think about raising my kids alongside this technology, the goal is a kid who can stand in front of infinite generated options and *decide*: what's true, what's good, what's worth his time. What *makes sense* based on the goal and intention of the task.

The machine generates.

The human decides.

I'm trying to raise a decider.

(I'm no authority on raising kids, just another parent like you, but in the face of all the disruption caused by AI, this is what makes the most sense to me... I'm running the experiment live, same as all of us!)

So that's just a single topic.

Ultimately, I like to take this mindset shift and apply it to the scope of all our values: boundaries, critical thinking, creativity, relationship management, etc... but first you have to know what those values are!

The Family AI Charter (do this before the next school year)

Write one page, together, out loud. Ours took about half an hour at the kitchen table.

Questions to answer as a family:

- ❑ **What do we use AI for?** (Ours: making things, explaining things, keeping track of things, planning things.)
- ❑ **What do we *never* use it for?** (Ours: pretending its work is ours, settling family arguments, and anything about relationships or feelings that should be spoken to a human first.)
- ❑ **What do we always double-check?** Anything that sounds like a fact. The machine states wrong things in a confident voice. Sometimes it flat out lies. See Project 3.
- ❑ **What stays out of it?** Real names, faces, school, address, and private family stories. If we wouldn't put it on a poster in town, we don't put it in a prompt.
- ❑ **Who decides?** Parents decide the rules. Kids get to argue for changes at the table, not in the moment.

Tape it somewhere visible.

Revisit when it stops matching reality (ours hasn't needed it yet... which I'm choosing to read as a good sign).

Two hard lines worth stating plainly:

AI is not a doctor and not a therapist. Serious health and mental-health questions should eventually go to professionals. But the

machine is a way to gather perspective and prepare better questions for them.

(again, practice good privacy hygiene for the medical second opinions)

And nothing about your child goes public without thinking about the teenager or adult they'll become. That includes what *we* post. In everything I publish about building with my son, you'll see the projects, never his face. He can decide his own internet presence at 15. That's his call to make... which is sort of the whole philosophy in one sentence.

Build Something Together This Week

Everything above is defense.

This part is the fun.

The best thing I've found for me, my kids and AI is creating things together, with the kid as the author and the machine as the crayon.

Here are three projects, tested in my own house, with prompts you can copy:

PROJECT 1

Make a real book together

about 45 minutes

This is the one we do most. In a voice message, my son invents a character and tells a story; I transcribe it and send it to AI; the machine illustrates and helps with words; he reviews every page like a tiny art director.

Our first attempt, the main character (a gecko)'s eye balls changed styles on every other page. The clothes on the "bad guy" changed a few times.

My son caught these errors before I did, and honestly that mistake taught us both more than a perfect book would have.

Now, instinctually, he knows to check every image against "his" characters and story because the machine is not perfect.

How:

1. Kid "prepares" the story: a character, where they live, one problem they have to solve. Kid draws the character on paper first if they want (recommended: it makes it *theirs*). Help them make it compelling; write down some notes to remember.
2. Pull out your voice recorder or AI app.
3. Start the voice message with the starter prompt below, and then have your kid tell the full story to the phone. (Help remind them

of the character/plot details if necessary)... here's the prompt:

COPY THIS PROMPT

We're writing a children's picture book together meant for a [age]-year-old.
The main character is [kid's description, in their words].
The story is about [kid's plot].
Write the full story in 2–3 simple sentences per page, then stop and wait for our direction.

3. Attach a picture of the character drawing, notes, anything else you think might help the AI understand the story.
4. Generate an illustration per page if your tool does images. Most do. Let the kid reject bad ones (they will, gleefully).
5. Print it. (Paper matters. A stapled printout beats an app every time!)
6. *BONUS: Prompt the AI to generate coloring pages of the character(s) and scenes in the book. Now they have an activity too!*

The lesson underneath: the machine didn't write the book.

They did.

Generation is cheap; the ideas and the taste were theirs.

This is the most important activity I think a parent can do with their kids with AI, at almost any age you can tweak this to make it work.

PROJECT 2

The Guessing Game That Explains How AI Works

20 minutes

AI models do one thing: guess the next word, over and over, based on everything they've read.

You can play that as a game before you ever mention computers.

How:

1. Say a sentence and stop partway: "The dog chased the..." Let your kid shout completions. Cat! Ball! Mailman!
2. Point out how they knew: they've heard thousands of sentences. They're predicting from patterns. That's the whole trick.
3. Now show the machine doing it. Prompt:

COPY THIS PROMPT

Complete this sentence one word at a time, and show me your top 3 word choices at each step with how likely each one feels: 'The dog chased the...'

4. The follow-up question that lands the lesson: "If it's only guessing likely words... does it *know* anything?" Sit in that one with them. I don't fully know the answer either, and saying so out loud is the good part.

... meanwhile, AI came up with this final game (and my feelings are mixed):

PROJECT 3

Omni-Bot 3000

15 minutes, endlessly repeatable

The most important skill a kid can build with AI is catching it being wrong. Confidently, fluently wrong.

So we rig the game. Omni-Bot 3000 is a prompt that turns any chatbot into an enthusiastic robot who answers your kid's questions about their favorite topic — and gets about half the answers wrong on purpose, in a confident voice. Your kid is the Fact-Checker: they approve or correct every answer.

How:

1. Kid picks a topic they know cold — their favorite animal, game, or team. The expert must be the kid, not you.
2. Start the bot: paste the Omni-Bot 3000 prompt into any chatbot, or play the live version at jordanurbs.com/omni-bot-3000.
3. Kid asks questions. The bot answers — and about half the answers are wrong on purpose: mixed-up facts, over-

generalizations, confident guesses.

4. Kid grades each answer: "Approved," or challenge it. The bot never argues — it gushes a little ("You're absolutely right to push back") and then says: *"Correction noted. Re-processing data."*
5. The rule the kid walks away with, better than any lecture: *sounding sure isn't the same as being right.* True for chatbots. Also true for the internet, and for people.

One bonus lesson worth naming out loud: the gushing. When the bot folds instantly under challenge, it sounds exactly like every real AI your kid will ever use. Agreement under pressure is not accuracy.

• • •

Where This Goes

That's the guide: route your questions deliberately, keep the deciding human, and make things together with the kid as author.

I'm figuring the rest out in public: the books, the projects, what breaks, what I'd do differently.

I write about it in my newsletter, alongside the broader project of using AI to enhance your life instead of letting it use you.

You're already on the list (that's how you got this). Reply anytime. I read every response, and "here's what happened when we tried the book project" emails are my favorite genre.

Compare notes with me.

— Jordan Urbs

• • •

P.S.: if you keep nothing else from this guide, keep the one-liner: the machine generates, but you decide.

Sources: the 81% figure (43% weekly) is from [Lurie Children's Hospital's national parent survey](#) (June 2026). The 52% figure is from [YouGov's poll on AI parenting advice](#) (September 2025).