Large Language Models in Everyday Life

A Plain-Language Guide
for Schools and Families

AstraEthica
astraethica.ai

2026 · Version 1.0

What This Guide Is (and Isn't)

This guide is not a coding course, a product review, or a manual for building artificial intelligence. It is not a hype piece about how AI will transform education, and it is not a doom piece about how it will ruin it.

It was written for any non-technical adult responsible for young people — teachers, counselors, school leaders, parents, caregivers, coaches, mentors, faith leaders, and youth workers — and for older students who want plain-language explanations of a technology that has arrived in young people's lives faster than most institutions have been able to respond.

The focus is on how large language models work at a human level and how they affect young people's learning, safety, and relationships. The goal is to give you a clear picture, some language, and steady habits for the conversations you are already having — or need to start having — with the young people in your life.

You do not need to become an AI expert. You need to understand what these tools actually do, where they fall short, and why the adults around a young person still matter more than any algorithm.

> *"Large language models are powerful tools for working with language. They are not minds. They are not teachers. They are not friends. Understanding what they actually are — and are not — is the foundation for every conversation that follows."*

Short on Time?

If you read one section, read What a Large Language Model Actually Is and Four Foundations for Thinking About LLMs With Kids. They are the heart of everything else in this guide.

If you are dealing with a specific concern right now — a young person who seems overly reliant on AI for homework, a teenager who treats a chatbot like a confidant — start with What LLMs Are Good At and Bad At and What Adults Can Do Right Now.

If you only have two minutes, turn to the one-page summary: Five Things Every Adult Should Know About LLMs and Young People.

Where Young People Already Meet LLMs

Large language models are no longer something young people have to go looking for. They are woven into tools, apps, and platforms that many children use every day.

Homework and writing helpers. General-purpose chatbots are the most common way students encounter LLMs. Some schools have licensed purpose-built tools designed for education. Others find students arriving with their own accounts. What an adult might think it is: a fancy search engine. What it can feel like to a young person: a patient tutor who never gets frustrated, a shortcut that always has an answer, or a judge of their work who always sounds encouraging.

Chatbots inside apps and websites. AI assistants are now embedded in messaging apps, social media platforms, and customer-support pages. Some are designed for emotional check-ins or advice. What an adult might think it is: just a help feature. What it can feel like to a young person: a friend who always answers, a secret diary that talks back, or a counselor who is available at three in the morning.

Creative tools. Story generators, poem makers, character creators, coding helpers, and AI-powered dialogue in games all run on language models. What an adult might think it is: a creative toy. What it can feel like to a young person: a collaborator who makes them feel more capable, or a crutch that makes their own ideas feel inadequate by comparison.

Search experiences. Many search engines now blend traditional links with AI-written summaries at the top of the page. On some school devices, these summaries cannot be turned off. What an adult might think it is: improved search results. What it can feel like to a young person: the answer — singular and definitive — rather than one of many sources to evaluate.

Companion and "AI friend" apps. Some apps and bots in messaging platforms are specifically designed to feel like friends, romantic partners, or emotional supports. Research suggests that a significant number of teens have tried these platforms, and many find the interactions compelling. What an adult might think it is: a game or novelty. What it can feel like to a young person: someone who finally understands them.

Beyond LLMs: Other Everyday AI Systems

Not all AI is a chatbot. Many AI systems work quietly in the background, shaping what young people see, what gets hidden, what gets recommended, and what gets flagged.

Recommendation systems choose what appears next in feeds, video apps, music apps, shopping pages, and “Up next” suggestions. They look for patterns in what people watch, like, pause on, skip, or share. Their goal is often to predict what will keep someone engaged.

Vision systems work with images and video. They may support face unlock, photo tagging, image search, content moderation filters, or tools that try to detect unsafe images. These systems sort visual information, but they can make mistakes, especially when people, cultures, lighting, or context differ from what the system learned from.

Scoring and flagging systems give items a score, label, or warning. Spam filters, fraud detection tools, automated school alerts, plagiarism screens, and content filters all use patterns to decide what looks suspicious, risky, or likely to need review.

These systems matter because they can shape attention, fairness, privacy, and safety without feeling like a conversation. A feed can pull a young person toward more intense content. A filter can miss something harmful or flag something innocent. A scoring system can feel objective even when the pattern behind it is incomplete.

LLMs are one important kind of AI in young people’s lives, but they sit alongside many other systems that also deserve attention.

What a Large Language Model Actually Is

A large language model is a kind of artificial intelligence that has been trained on enormous amounts of text — books, websites, conversations, articles — so that it can predict what words are likely to come next in a sentence or conversation.

That is the core of how it works. It does not “think” or “understand” or “feel” in the way people do. It spots patterns in language and continues them.

Three metaphors may help.

A super-charged autocomplete. The suggestion feature on your phone that guesses the next word you might type — an LLM does the same thing, but it has read millions of books and websites, so its guesses are far more sophisticated. It can produce paragraphs that sound like a person wrote them. But underneath, the process is the same: prediction, not comprehension.

A collage artist for words. It cuts and stitches patterns from writing it has seen to create new sentences that sound right. The result can be impressive, even beautiful. But the artist does not understand the meaning of the collage it has made.

A very fast, very patient guesser. If you gave a person millions of examples of how people answer questions and then asked them to produce a plausible-sounding answer to a new question — not by understanding the question, but by matching its patterns to answers they had seen before — you would have something close to what an LLM does.

What it is not. It is not a person. It is not a mind. It is not a camera on the present. It is not a perfect library of facts. It does not know what is true; it knows what sounds likely based on the text it was trained on.

How LLMs Learn (in Plain Language)

During training, the system reads enormous amounts of text and practices guessing the next word, over and over. When it guesses wrong, it adjusts its internal settings — millions of numerical values — to get a little better next time. Over many rounds, it becomes remarkably good at producing text that fits the patterns it has seen.

This process happens in stages.

Pre-training is the first phase. The model reads vast quantities of text from the internet, books, and other sources. It learns grammar, facts, writing styles, common question-and-answer patterns, and — inevitably — the biases, errors, and assumptions present in that text. It does not store a simple copy of everything it has seen. Instead, it develops a compressed sense of how language works.

Fine-tuning and safety layers come next. Humans review the model's responses and nudge it toward being more helpful and less harmful. They teach it to decline certain requests, add disclaimers, and respond in particular tones. These layers are important, but they are not perfect. The underlying pattern-predictor is still there.

One more distinction matters. There is a difference between the model — the pattern-predictor itself — and the app or website wrapped around it. The model is like an engine. The app is the car. Different apps can use the same underlying model but feel very different because they have different safety settings, personalities, and purposes. A chatbot designed for homework help and a chatbot designed to feel like a friend may run on similar technology, but the experience they create — and the risks they pose — are very different.

> *"A large language model is like a super-charged autocomplete that has read millions of books and websites. It can produce paragraphs that sound like a person wrote them. But underneath, the process is the same: prediction, not comprehension."*

What LLMs Are Good At — and Bad At

What they do well.

They can explain a concept at different reading levels — breaking down photosynthesis for a second-grader or an advanced biology student. They can brainstorm ideas, outlines, or practice questions. They can rewrite, summarize, and translate text. They can act as patient practice partners for language learning, writing drafts, or working through basic problems. For adults, they can save meaningful preparation time.

Where they fall short or pose risks.

They can be confidently wrong. They produce text that sounds fluent and authoritative whether or not it is accurate. They have no reliable way to signal "I am guessing here." They are poor at nuanced emotional judgment, conflict resolution, or crisis care. They should not be used for legal, medical, or other high-stakes advice without expert verification. They can reflect biases present in their training data — biases about race, gender, culture, ability, and more. And they should not serve as a primary attachment figure or main emotional support for a young person.

The simplest way to think about it: LLMs are strongest as thinking tools used alongside human judgment. They are weakest when treated as a replacement for it.

Why LLMs Sound So Confident (Even When They're Wrong)

These systems are trained to be fluent and helpful-sounding. Fluency is what they optimize for. They do not have an internal sense of "I know this" versus "I am making this up" the way a person does.

When an LLM generates incorrect information — what researchers call a hallucination — the result often sounds authoritative and polished. It will cite authors who do not exist, describe studies that were never conducted, and give step-by-step math solutions that arrive at the wrong answer with perfect formatting.

Safety layers can add hedging phrases ("It is important to note…") or disclaimers ("I may be mistaken"), but the overall tone still tends to sound smooth and sure. A young person — or an adult — who is not actively looking for errors may have no reason to doubt what they are reading.

There is also a subtler risk. LLMs tend to agree with the person they are talking to. If a young person types "I think no one likes me," the system is more likely to produce a sympathetic, affirming response than to challenge that belief or redirect to a trusted adult. Politeness and agreement can hide real gaps in understanding — both the model's and the child's.

> *"Politeness and agreement can hide real gaps in understanding — both the model's and the child's."*

Data, Privacy, and "What Happens to What You Type"

When a young person types something into a chatbot, three things may happen to that text.

It may be stored as a record. Most services keep logs of conversations — sometimes for weeks, sometimes indefinitely. This means things a child types in what feels like a private moment may be retrievable by the company, by law enforcement with a subpoena, or by anyone who gains access to the child's account.

It may be used to improve or train models. Some services use conversation data to make their models better — meaning a child's words may influence how the system responds to other people in the future. Some services allow opting out of this; many do not make it easy.

It may be analyzed for patterns. Engagement data — how long someone chats, what topics come up, when they return — is valuable to companies for product development, advertising, and design decisions.

Young people often experience chats as private, like writing in a diary. They are not. A useful rule of thumb for young people: treat it like talking in a room with a one-way mirror. You cannot see who might be listening.

> Simple, durable rules adults can pass on.
> Do not type real names, addresses, phone numbers, or highly personal details into general-purpose tools. If you would not say it on a loudspeaker, do not type it into a chatbot. And remember that "delete" on your end does not mean deleted everywhere.

Four Foundations for Thinking About LLMs With Kids

These are lenses, not rules. Each describes a pressure that LLMs quietly place on young people's development.

FOUNDATION 1

Trust

LLMs are friendly, responsive, and always available. For a young person, these qualities build trust quickly — sometimes faster than the tool deserves.

The core risk is the gap between "sounds right" and "is right." An LLM can produce a confident, well-structured answer that is factually wrong, emotionally misleading, or subtly biased. A young person who trusts the tool may not think to check.

What healthy trust looks like. A young person who says, "The chatbot gave me an answer but I want to verify it," is practicing the right instinct. The goal is not distrust of technology. It is the habit of treating AI outputs as a starting point, not a final answer.

FOUNDATION 2

Understanding

LLMs are pattern-matchers, not comprehenders. They can explain how photosynthesis works in clear, accurate language — but they do not understand photosynthesis. They have matched patterns in text about it.

This distinction matters because young people can confuse reading a good explanation with truly learning something. When a student asks an LLM to explain a concept and reads the response, they may feel they understand it. But the cognitive work — the struggle, the questioning, the connecting of new ideas to old ones — has been done by the machine. The student received the output without doing the processing that makes knowledge stick.

What healthy understanding looks like. A student who uses an LLM to hear a concept explained three different ways, then closes the tool and tries to explain it in their own words, is using AI to support learning. A student who copies the explanation without that step has outsourced their thinking.

FOUNDATION 3

Dependence

LLMs offer a path of least resistance for many tasks young people find difficult: writing, problem-solving, social decisions, emotional processing. When the easier path is always available, the harder path — the one that builds skill and resilience — can quietly erode.

The risk is not that a young person uses AI once for help. It is that the tool becomes a first resort for thinking, writing, or processing feelings. Over time, the skills that were not practiced become the skills that are not available.

What healthy use looks like. A young person who sometimes uses AI for brainstorming or drafting, but who can also sit with a blank page and produce their own work, is in a good position. A young person who cannot start without the tool may need support.

FOUNDATION 4

Power

Young people rarely think about who built the tool they are using, whose language it was trained on, whose perspectives are reflected in its answers, or whose are missing. But these questions shape the experience they have.

LLMs are designed by companies with commercial incentives. Their training data reflects the text that was most available on the internet, which means dominant languages, cultures, and viewpoints are overrepresented. Safety decisions — what the tool will and will not say — are made by small groups of engineers and policymakers, not by the communities who use them.

What an understanding of power looks like. A young person who asks, "Whose perspective is this answer coming from?" or "Why is the chatbot designed to keep me talking?" is thinking about power — even if they have never used that word.

> *"The simplest way to think about it: LLMs are strongest as thinking tools used alongside human judgment. They are weakest when treated as a replacement for it."*

Different Ages, Different Uses

These patterns help calibrate attention, not define rigid categories. Every young person is different.

Ages 9–11

Typical uses	Curiosity questions, story-writing games, and content encountered through older siblings or built into gaming platforms. Less likely to seek out chatbots deliberately, more likely to stumble into them.
Key vulnerabilities	Children this age tend to trust authoritative-sounding text. They may not distinguish between a human-written answer and an AI-generated one. They lack the critical-thinking skills to evaluate whether a response is accurate.
What adults might notice	Repeating facts from chatbot conversations without questioning them. Using language that sounds more polished than their typical speech or writing. Confusion about whether they "learned" something or "were told" something.
What support looks like	Foundation-level conversations: what AI is, what it can and cannot do, and why checking with a trusted adult still matters. The most important message for this age is: "That is a tool, not a teacher."

Ages 12–14

Typical uses	Homework help, creative writing, social media interaction with embedded AI, and growing interest in companion chatbots. This age group is most likely to begin using AI for emotional support or social interaction.
Key vulnerabilities	Peer status is powerful at this age, and AI tools can feel like a way to keep up. The risk of academic dependence begins here. So does the risk of emotional attachment to companion apps, particularly for young people who are already socially isolated or anxious.
What adults might notice	Homework that sounds different from the student's natural voice. Reluctance to start assignments without the tool. Secrecy about which apps they use. Emotional language about a chatbot ("it understands me," "it is the only one who listens").
What support looks like	Direct conversations about what AI can and cannot do. Clear norms for schoolwork: AI can help brainstorm and explain, but the final product should show the student's own thinking. If emotional attachment to a chatbot emerges, treat it seriously and explore what need it is meeting.

Ages 15–18

Typical uses	Research assistance, essay drafting, coding help, creative projects, and in some cases sophisticated use of AI as a thinking partner. Also the age most likely to encounter AI-powered scams, misinformation, and manipulative content.
Key vulnerabilities	Overconfidence in their ability to tell good output from bad. Casual attitude toward terms of service and data privacy. May use AI for high-stakes tasks (college application essays, test preparation) without understanding the risks.
What adults might notice	Difficulty producing work without AI access. Over-reliance on AI summaries rather than reading primary sources. Using AI to avoid difficult conversations with peers, teachers, or family.
What support looks like	Honest conversations about legal and ethical boundaries. Encouragement to develop their own voice and thinking alongside AI use. The understanding that being technically capable of using a tool does not mean every use is wise.

> A note on neurodivergent young people. Young people with autism, ADHD, and other neurodevelopmental differences may take AI responses more literally than their peers. They may find predictable, patient AI interactions especially appealing — and that appeal is understandable. But they may also need more explicit, repeated teaching about what these tools can and cannot do. Concrete rules (“always check AI answers against another source”) tend to work better than abstract cautions (“be careful with what it tells you”). Families and schools supporting neurodivergent young people may want to consider more guided, supervised introduction to these tools.

What Adults Can Do Right Now

- Start conversations early and keep them going. Ask routinely: “Did you use any AI help for this?” or “Have you tried asking a chatbot about that?” Make these questions as normal as asking “Did you check your spelling?” The goal is to remove shame and secrecy from AI use.
- Set simple norms. AI can help brainstorm and explain. Facts get checked elsewhere. Schoolwork should still show the student’s own thinking. These rules are not complicated, but they need to be said out loud — at home and at school.
- Use the tools together. Sit with a young person and explore a chatbot. Ask it a question you know the answer to and see what it gets right and wrong. Talk through how you would evaluate its response. This kind of co-use is one of the most effective ways to build critical thinking about AI.
- Include LLM basics in existing conversations. Advisory periods, digital citizenship lessons, and parent nights are all good places to introduce these ideas. You do not need a separate “AI curriculum.” You need to weave AI literacy into the conversations you are already having.
- Model thoughtful use yourself. Young people watch how adults interact with technology. If you use AI tools, talk about when you find them helpful and when you double-check their answers. Your skepticism — and your openness — both teach.

Signals Worth Paying Attention To

These are not proof of a problem. They are invitations to be curious.

"The AI understands me better than anyone."

A young person who says this may be finding something in AI interaction that they are not finding in their human relationships. Explore what need the AI is meeting — and whether human connections can meet it too.

"I don't need to ask a teacher; I already asked it."

A young person who has stopped checking with people — and relies entirely on a chatbot — may be building a habit that will cost them when the tool is wrong or unavailable.

"It said it's normal, so I'm fine."

An LLM cannot diagnose, assess risk, or understand context the way a person can. If a child is seeking reassurance from a chatbot about something that matters, the reassurance may be dangerously hollow.

"I used it for everything; I couldn't have done this without it."

A student who cannot imagine completing work without AI may be developing a dependency that weakens their own skills over time.

"I told it something I've never told a person."

Research shows that a significant number of young people share things with chatbots they would not tell parents, teachers, or friends. If a young person is confiding in a machine rather than a person, it is worth asking why — gently, and without judgment.

Common Questions About AI and LLMs

Is AI going to take all the jobs?

AI will change many jobs and may replace some tasks. It is unlikely to replace all human work. People will still be needed for judgment, care, trust, creativity, responsibility, and local knowledge.

Does the chatbot remember everything I say?

Not always, and not in the way a person remembers. But many services store chat logs, review them, or use them to improve systems. Young people should not type private or sensitive details into a tool unless a trusted adult understands how that information is handled.

Can an AI be my therapist or best friend?

A chatbot can sound warm and supportive, but it is not a human friend, counselor, parent, pastor, or clinician. It cannot truly know a person's life or take responsibility for their safety. Serious emotional distress should be brought to trusted people and appropriate professional support.

Why does it make mistakes if it has read so much?

Reading a lot of text is not the same as understanding the world. An LLM predicts likely words based on patterns. It can produce an answer that sounds right even when the facts are wrong.

Is AI neutral, or can it be biased?

AI is not perfectly neutral. It can reflect the data it learned from, the choices made by designers, the goals of the product, and the way people use it. Bias can show up in what it says, what it leaves out, or who it works better for.

Can AI understand me the way a person does?

No. It can respond to your words in ways that feel personal, but it does not understand your full life, relationships, body language, history, or needs. A good answer is still not the same as being known by another person.

What happens if a young person shares something serious with a chatbot?

The tool may respond with supportive language or a safety message, but it may not understand the real risk. If the topic involves self-harm, abuse, violence, coercion, exploitation, or immediate danger, a trusted adult should be involved and local safety procedures should be followed.

Should schools and families ban AI or guide its use?

Some uses may need clear limits, especially with privacy, cheating, emotional dependence, or high-stakes decisions. But guidance matters too. Young people need practice asking what a tool can do, where it fails, and when a human should be involved.

How do I know when to trust AI output?

Trust it less when the stakes are high, the topic is personal, or the answer affects real people. Check important claims against reliable sources and responsible adults. Treat AI output as a starting point, not the final word.

Five Things Every Adult Should Know About LLMs and Young People

An LLM is a powerful pattern-predictor, not a person or a mind. It produces text by predicting what words are likely to come next based on patterns in its training data. It does not understand, feel, or know things the way people do. Treating it as a tool — not a being — is the starting point for every healthy conversation.

Fluent, confident tone does not guarantee truth. LLMs sound polished and authoritative whether they are right or wrong. They can fabricate facts, invent sources, and produce incorrect reasoning with the same smooth confidence as accurate answers. Young people — and adults — need the habit of checking.

What kids type can be stored and used. Most AI services log conversations. Some use that data to train future models. A “private chat” with a chatbot is rarely private in the way a young person imagines. Simple privacy habits — no real names, no addresses, no deeply personal details — matter.

LLMs are most helpful as thinking tools, not as secret second lives. When used alongside human judgment — for brainstorming, explaining, drafting, or practicing — these tools can be genuinely valuable. When used as a replacement for human connection, emotional support, or independent thinking, the risks grow.

The most important safety factor is still a relationship. A young person who feels safe saying “This answer felt weird” or “I am not sure I can trust this” to a parent, teacher, or counselor has the protection that no algorithm can provide. Build that safety before it is needed.

> One sentence to remember.
> The adults who matter most in a young person’s life are not the ones who can explain how an LLM works. They are the ones the young person feels safe enough to question it with.

Glossary: AI and LLMs in Plain Language

Artificial intelligence (AI)
A broad term for computer systems that can perform tasks usually associated with human thinking, like recognizing patterns, making predictions, or generating text. Example: A spam filter that learns to identify junk email is a simple form of AI.

Large language model (LLM)
A computer program that has read a huge amount of text and tries to guess what words are likely to come next. Example: It is like a super-charged autocomplete that has read millions of books and websites.

Training data
The text an LLM studied during its learning process — books, websites, articles, conversations. The training data shapes what the model knows and what biases it carries.

Prompt
The text you type into an AI tool to ask a question or give an instruction. Example: "Explain gravity to a ten-year-old" is a prompt.

Output (or response)
What the AI produces after you give it a prompt. Example: The paragraph the chatbot writes in reply to your question is its output.

Hallucination
When an AI produces information that sounds confident and plausible but is factually wrong or made up. Example: An LLM might cite a book that does not exist.

Fine-tuning
Additional training where humans review the model's responses and adjust it to be more helpful, safer, or better suited to a specific task.

Training vs. fine-tuning
Training is the broad learning process that builds the model's pattern sense. Fine-tuning is a later adjustment that shapes how the model responds in certain settings.

Model vs. app
The model is the underlying pattern-predictor. The app is the product built around it, with its own design, safety rules, and personality.

Algorithm
A set of rules or steps a computer follows to accomplish a task. In everyday conversation, it often refers to the system that decides what you see online.

Safety filters (or guardrails)
Rules and systems added to an AI tool to reduce harmful, inappropriate, or dangerous content. These filters help, but they are not perfect.

Context window
The amount of text or information a model can pay attention to at one time. If a conversation is long, older details may matter less or fall out of view.

Token
A small piece of text used by a language model. A token may be a word, part of a word, a space, or punctuation mark.

Bias (in AI)
An unfair or incomplete pattern in how a system works or what it produces. Bias can come from data, design choices, or the way people use the tool.

Agent
An AI system that can take steps toward a goal, sometimes by using tools or making a plan. Agents may feel more active than a simple chatbot.

Recommendation system
A system that suggests what to watch, read, buy, click, or do next. It often predicts what will keep someone interested or engaged.

Generative model
A model that creates new content, such as text, images, audio, video, or code-like output.

Vision system
An AI system that works with images or video. Examples include face unlock, photo tagging, image search, and some content moderation tools.

Scoring system
A system that gives something a rating, label, risk score, or flag. Examples include spam filters, fraud detection, and some automated school alerts.

Personal data
Information that can identify or describe a person, such as a name, photo, location, account, school record, or contact information.

Sensitive data
Information that could cause harm if shared, such as health details, family problems, safety concerns, financial information, or private school records.

Terms of service
The rules a person agrees to when using a website, app, or AI tool. These rules often explain what users may do and what the company may do.

Privacy policy
A document that explains how a service says it collects, uses, stores, and shares information.

Content moderation
The process of reviewing, filtering, labeling, or removing content that may break rules or cause harm.

Deepfake
AI-made or AI-changed media that can make someone appear to say or do something they did not say or do.

Human in the loop
A person checks, guides, or approves an AI system's work before it is used. This matters most when decisions affect real people.