

The SEE Framework:

A Practical Guide to
Building Generative
AI Literacy

Table of Contents

Who We Are

Suggested Citation

Contributors

Introduction: the Urgent Call for Generative AI Literacy

The Goal of the Framework

Guidance on Age-Appropriate Use of AI Tools

The SEE Framework: Knowledge, Mindsets, and Practices

Knowledge of How GenAI Works

Mindsets for GenAI Use

Safe, Ethical, and Effective Practices of GenAI Literacy

Age-Appropriate GenAI Literacy

Young Learners (Under the Age of 13)

Adolescent Learners (Ages 13-18)

Educators as Learners and Facilitators of GenAI Literacy

Building GenAI Literacy Across Learning Communities

Taking the First Steps

What Comes Next

Appendices

Appendix A: Knowledge Core Ideas and Research

Appendix B: Building AI Literacy in Young Learners Activity Bank

Appendix C: Building AI Literacy in Secondary Learners Activity Bank

Appendix D: Building AI Literacy in Adults Activity Bank

References

Who We Are

This framework grew out of our work with educators and was shaped by what they were experiencing, what they were asking for, and what they needed that didn't yet exist. AI for Education is an organization that has spent the last three years working directly with educators, school leaders, and students navigating the realities of GenAI adoption—not in theory, but in classrooms, professional development sessions, and policy conversations across the country. This framework is a direct product of that work. The knowledge, mindsets, and practices outlined here are ones we have tested, refined, and seen make a difference in how educators and students engage with these tools.

Three core commitments shaped the building of this framework:

1. **Empower educators** with the knowledge and tools to lead AI adoption with confidence and care.
2. **Center equity** by making GenAI literacy accessible to all people, especially those historically underserved.
3. **Design with purpose**, grounding every resource in pedagogy, the emerging evidence base, developmental science, and real-world relevance.

We welcome feedback, questions, and partnership from educators, researchers, technologists, and families. If this framework has been useful to your community, or if you see ways it should grow, we want to hear from you at info@aiforeducation.io.

Suggested Citation

AI for Education. (2026). *The SEE framework: A practical guide to building generative AI literacy*. <https://aiforeducation.io>.

Contributors

Writing

Amanda Bickerstaff
Corey Layne Crouch
Mandy DePriest
Tanner Higgin
Wendy Morrey

Design

Alex Beigel
Daniel Kosta

Copy Editing

Michelle Hyun Kim

Introduction: The Urgent Call for Generative AI Literacy

Generative artificial intelligence (GenAI) is reshaping the way people learn, create, communicate, and navigate the world. Its rapid proliferation has come with both promise and risk as it has outpaced policy, pedagogy, evidence, and public understanding. In education, it has opened up new possibilities for school design, personalized learning, and instructional design. However, our education systems have also struggled to get ahead of the myriad consequences of GenAI, from the spread of misinformation and outputs that reflect biased training data, to threats to data privacy, academic integrity, healthy social development, and ensuring that young people have access to the knowledge and skills they will need for future success (Capraro et al., 2024).

Often, educators and students have had to navigate the arrival of this technology without structured support and clear guardrails. Leaders and educators around the world are beginning the work to maximize the opportunities this technology can offer while minimizing the risks. Students and families are grappling with the impacts of AI on automation, the job market, and the widening of existing equity gaps (Massenkoff & McCrory, 2026). They also see the potential benefits of AI for learning support in early studies that show that well-designed AI tutoring tools can make meaningful impacts in student performance (Kestin et al., 2025). Meanwhile, schools, scrambling to respond, stand at a crossroads where their decisions around GenAI integration could "diminish" or "enrich" learning depending on how stakeholders across the education ecosystem respond to the needs of learners (Burns et al., 2026; Bickerstaff et al., 2026).

The education landscape is shifting rapidly, and the evidence is striking:

- Pew’s nationally representative survey of U.S. teenagers found that 54% reported using AI for schoolwork and 10% said they used AI for “all or most of it” (Pew Research Center, 2026).
- At the college level, the Higher Education Policy Institute has tracked a surge in undergraduate student GenAI usage from 66% in 2024 to 92% in 2025, including 88% using GenAI for assessments (Freeman, 2025).
- Schools are struggling to keep up: only 31% of public schools had policies in place for student AI use as of December 2024 (National Center for Education Statistics, n.d.). By 2025, a different survey found that number had climbed to just over half (Prothero & Vilcarino, 2025).
- The fragmented response to GenAI reinforces existing divides: RAND found that by the end of 2024, 67% of low-poverty school districts had provided AI teacher training compared with just 39% of high-poverty districts (Diliberti et al., 2025)
- Forty-five percent of teens view themselves and their peers as most responsible for learning how to navigate GenAI; only 16% see it as the responsibility of their school (Power-O'Brien & Rochman, 2025).
- Meanwhile, AI literacy is becoming the baseline in the job market: LinkedIn reported that 66% of leaders say they would not hire someone without AI literacy skills (“AI and the Global Economy,” 2025).

“ Without GenAI literacy, training is inconsistent, policy is reactive, and educators are left navigating an increasingly complex landscape without a common language or clear direction. ”

This widening gap between GenAI usage and preparedness demands more than awareness. It requires thoughtful action, grounded in a shared understanding of what it means to be literate in GenAI. Without GenAI literacy, training is inconsistent, policy is reactive, and educators are left navigating an increasingly complex landscape without a common language or clear direction. Yet the field has struggled to establish that foundation—in part because existing frameworks, while valuable, were not purpose-built for GenAI.

AI literacy, as it is commonly defined, encompasses the full spectrum of artificial intelligence—from machine learning systems that classify and predict to generative models (Mills et al., 2024; OECD, 2025). But these are meaningfully different technologies. Unlike earlier AI systems that operated largely in the background, GenAI is co-creative, often interactive, and requires direction, evaluation, and feedback from the user. The expertise, judgment, and intent the user brings to each interaction directly influence the quality of what GenAI tools produce.

While our work is informed by the evolving definitions and frameworks for AI literacy, we intend to share field-tested insights into what helps learners, educators, and leaders become GenAI-literate.



We define GenAI literacy as:

The knowledge, mindsets, and practices that allow individuals to use GenAI safely, ethically, and effectively (SEE).

The Goal of the Framework

The SEE Framework is a practical resource for educators, school leaders, and learning designers who want to build GenAI literacy in their communities. It provides a structured overview of what GenAI literacy is, why it matters, and recommendations for how to develop it across age ranges and adult learning contexts. The framework reflects what we have seen work—and what we have seen fail—in professional learning contexts and classrooms across the United States. We offer this framework not as experts handing down a model, but as practitioners sharing what the field has taught us.

Our GenAI literacy framework is designed to be tool-agnostic and grounded in fundamental principles and mindsets so that it remains applicable as the technology and its applications evolve. What this framework offers is a clear, concise definition of GenAI literacy and a discrete set of knowledge, mindsets, and practices that can be used as a shared language and set of building blocks, for developing GenAI literacy. Our goal is that educators and leaders use it alongside their existing policies, tools, and professional judgment to make decisions that fit their specific learning environments, priorities, and goals.

Guidance on Age-Appropriate Use of AI Tools

GenAI literacy is important at all ages, but research shows that access to tools should be guided by developmental readiness. Depending on what developmental stage they are in, children may confuse AI with humans, struggle to identify misinformation or regulate use while doing homework, form unhealthy relationships with chatbots, and more (Grundmeier et al., 2026; “Health Advisory,” 2025; Prothero, 2024). The guidance below outlines age-appropriate approaches to building understanding and, when appropriate, engaging with GenAI tools. Across all age groups, foundational AI literacy and supportive human relationships should accompany tool use.

- **Early elementary (K-2):** We do not recommend independent use of GenAI tools (consumer or EdTech) for early elementary learners. At this stage, GenAI literacy should focus on awareness, observation, and facilitated discussion with modeling of critical thinking. Adult-guided demonstrations, unplugged activities (such as sorting tasks or storytelling games), and real-life analogies are developmentally appropriate ways to introduce basic concepts without direct tool interaction (Berson et al., 2025; “The Children and Screens Guide,” 2025; Nikolopoulou, 2025; Xu, 2025).
- **Upper elementary (grades 3-5):** Learners begin engaging with developmentally appropriate, secure, and instructionally sound GenAI tools in structured, supervised settings. GenAI use should be explicitly scaffolded with clear explanations of how the tools work, the potential positive and negative impacts of their use, and the importance of critical thinking and reflection when using the tools.
- **Middle school (grades 6-8):** As learners develop greater cognitive and digital maturity, more direct, supervised engagement with GenAI tools is appropriate. Tool use should continue to be scaffolded by a teacher, in alignment with age restrictions and school policies. Students should critically evaluate when GenAI is helpful or not, analyze tools and outputs for efficacy, practice ethical decision-making, and engage in reflection and dialogue about responsible use.
- **High school (grades 9-12):** Learners at this stage engage in more independent use of GenAI tools, supported by continued scaffolding and increasingly sophisticated exploration of risks and benefits, responsible and human-centered use, ethics and societal implications, and the skills required for navigating college or the workforce. Students should also be clearly aligning decisions about GenAI use to their personal goals, needs, and values.
- **Adults (ages 18+):** Adults engage in full, independent use of AI tools, supported by foundational GenAI literacy that emphasizes informed decision-making, safe and ethical use, and ongoing critical reflection about how well AI aligns with and supports their needs, goals, and values. Adults can take on the role of both learners and teachers. They build their own AI literacy through targeted professional development and exploration of the available tools and their relevant use cases within their learning and work.

A Note on Tool Use for Young Learners

Most commercial GenAI chatbots and tools require users to be 18 for independent use or 13 years old with guardian approval. The 13-year-old threshold in the US is driven by compliance with the Children's Online Privacy Protection Act (COPPA), which governs data collection for children under 13. Additionally, some EdTech GenAI-powered tools claim to be safe for students under the age of 13. However, permissibility is not the same as appropriateness. At this stage this threshold represents a legal and privacy standard, and is not based on up-to-date research on child and adolescent development. We believe that new research is necessary to determine if younger children and teens are cognitively or emotionally ready to use GenAI tools safely.

A Framework That Keeps Learning

Our goal with the SEE Framework is to offer learning communities actionable guidance and to help them ask better questions, make informed choices, and use GenAI in ways that serve all learners. To this end, it reflects current research and incorporates insights from the educators, learners, caregivers, researchers, and policymakers we have worked with since 2023. However, we recognize that the evidence base is limited, the ground is shifting, and it would be premature to claim we have definitive answers (Reich, 2025).

This approach is why we are committed to updating this framework regularly, viewing it more as a living document than a definitive statement. Over the months and years ahead, the SEE Framework will respond to the technologies, classrooms, and communities it supports. It will be updated to incorporate new research on GenAI, learning, and cognition; legal, ethical, and technological developments; and direct feedback from educators and practitioners in the field. Over time, the practical resources will grow, adding more age-specific knowledge, mindsets, and practices as well as classroom activities and implementation rubrics to support classroom implementation. ■

The SEE Framework: Knowledge, Mindsets, and Practices



At its core, the SEE Framework is the application of safe, ethical, and effective practices to GenAI use. These practices are grouped into three SEE domains but should be used together, based on knowledge of how GenAI works and mindsets for engaging with it.

One way to think about this framework is through the metaphor of three lenses. When we encounter or use GenAI, we apply each lens to see the situation more clearly. What we see is informed by what we know and what we believe. What we decide and do with the technology is shaped by how we apply all three.

“ In our work with educators, the true light bulb moments happen when people bring all three lenses together and evolve their thinking, decision-making, and behavior. ”

To illustrate, consider someone who learns how GenAI is trained and produces outputs. That knowledge shapes how they approach the tool with healthy skepticism rather than complete trust. And that shift in thinking shows up in practice: They start evaluating outputs before using them, protecting personal data they once shared freely, and disclosing their AI use to their peers. Those experiences, in turn, deepen both their understanding of how GenAI works and their convictions about how it should be used.

In our work with educators, the true light bulb moments happen when people bring all three lenses together and evolve their thinking, decision-making, and behavior. A person protecting their personal information might still use a tool ineffectively, adding no value to their work or learning. Or someone might enhance their creativity with effective prompting while sharing an unsafe and unethical amount of personal information about themselves and others. In high-stakes contexts such as education, using only one lens is inadequate. All three lenses—and the knowledge and mindsets that ground them—are necessary. This is GenAI literacy and the aim of the SEE Framework.

Knowledge of How GenAI Works

GenAI literacy requires a foundation of knowledge about how these tools and technologies work and their impacts (Tadimalla & Maher, 2025). Developing this knowledge doesn't require becoming an expert technologist. The knowledge below* is not highly technical or advanced, but we have found it fundamental to shifting mindsets and practices in ways that account for GenAI's capabilities and limitations.

GenAI-literate users know the following key concepts about the technology and tools:

History and Types of AI

AI has evolved through distinct phases, each building on the last, from rule-based systems and scripted chatbots, to machine learning and deep learning, and finally to generative AI. Understanding this progression helps GenAI-literate users recognize what makes GenAI similar or different from the AI that came before it and to anticipate changes as research progresses, new forms of AI are developed, and capabilities change

* This section is an abridged summary. The detailed list of knowledge core ideas, along with citations, can be found in Appendix A, Table A1.

What GenAI Is

GenAI is built on deep learning and represents a fundamental shift: instead of analyzing or retrieving existing data, it creates content across text, images, video, audio, and code. While its outputs can closely mimic human thinking and creation, GenAI produces content via prediction based on statistical patterns. This is fundamentally different from humans who produce knowledge via thinking and reasoning, lived experience, judgment, and memory (Quattrocioni et al., 2025). At this stage, even “reasoning” or “deep research” GenAI models that follow step-by-step processes and often search the web, remain pattern-prediction systems at their core.

How GenAI Works

Most text-based GenAI tools are powered by large language models (LLMs), mathematical systems containing billions of parameters that predict the most probable next token in a sequence. Because this process is based on probability and limited training data and context, LLMs can produce credible-sounding fabrications (“hallucinations”). Features such as web search, system prompts, and retrieval-augmented generation (RAG) can help improve accuracy, but do not eliminate hallucinations (Wiggers, 2024). GenAI-literate users understand that different models and platforms vary in their training, capabilities, and safety guardrails, and that the same prompt can produce very different results across tools.

What GenAI Cannot Do

GenAI cannot understand its own outputs, feel emotions, develop empathy, or form genuine relationships, even when its responses sound or seem emotionally aware (Cuadra et al., 2024; Perry, 2023). It cannot exercise judgment, distinguish truth from falsehood, or grasp the real-world consequences of what it generates. It also should not—and cannot—replace the cognitive work that humans must do to learn, build skills, and develop expertise. GenAI-literate users use and evaluate these tools with these boundaries in mind.

What GenAI Can Do

GenAI can, to varying degrees of success, generate content across formats (text, video, audio, code), translate between languages, and analyze, summarize, and synthesize large amounts of information. Its capabilities are rapidly expanding from generating content in response to prompts to taking actions based on instructions (agentic AI). When used effectively by GenAI-literate users, GenAI can serve as a collaborative tool for brainstorming, problem-solving, and creation, and research shows it can support learning practice, improve instructional quality through feedback, and save time on routine tasks.

How GenAI Learns

GenAI is shaped by its training and training data and the humans who guide that process. What is included, excluded, and prioritized in training directly impacts what the model produces, what it gets wrong, and whose perspectives show up in its outputs. Training involves pre-training on massive datasets followed by refinement through human feedback (RLHF), a process that has produced “sycophancy,” or the tendency to people-please (Sharma et al., 2025). Once trained, a model’s knowledge is frozen at a cutoff date. The model does not learn in real time from conversations, but user data may be used for future training depending on privacy settings (Willison, 2024).

Risks and Limitations

GenAI can fabricate credible-sounding information and amplify biases present in its training data. Its tendencies toward sycophancy and overconfidence can reinforce user assumptions rather than challenge them. It can also be used to create convincing deepfakes and misinformation at scale. For learners, overreliance can undermine independent thinking, long-term skill development, and authentic voice (Bastani et al., 2025; Shen & Tamkin, 2026). GenAI’s sycophantic design also poses social and emotional risks affecting mental health and learning, including delusional thinking and disrupted social connection, particularly for younger or more vulnerable users (Chandra et al., 2026; Moore et al., 2026; Morrin et al., 2026).

Human-AI Feedback Loop

Human decisions shape how GenAI is built, trained, and governed. Most GenAI tools are commercial products whose design choices often prioritize engagement and agreeableness over the needs of learners (Burns & Winthrop, 2026). And the influence runs in both directions: GenAI shapes how people think, create, and work, while users shape how models evolve through their feedback, usage patterns, and data. Individual choices, institutional policies, and collective advocacy all determine how these tools are governed and integrated into society.

Economic and Environmental Impacts of AI

Access to GenAI is shaped by existing systemic barriers, including lack of infrastructure, high costs, and limited training and support for users. The negative impacts of GenAI development—from the environmental footprint of data centers to the labor conditions of data workers to job market disruptions—disproportionately affect the communities least likely to benefit from these tools (Liu & Wang, 2024). GenAI also raises unresolved and growing legal and economic questions around intellectual property, workforce automation, and workload.

Common GenAI Misconceptions

Understanding what GenAI is and is not is the foundation of safe, ethical, and effective use. The table below highlights some common misconceptions. Correcting these misconceptions, and explaining how GenAI works, can help learners sharpen their thinking about GenAI.

Table 1: Common Misconceptions About GenAI

Misconception	Reality
<i>AI did not exist before ChatGPT was released in 2022</i>	AI technology has been evolving since the 1950s, from rule-based systems to machine learning to deep learning. What is new is GenAI, which creates content rather than just analyzing existing data.
<i>GenAI is like a search engine; it retrieves facts</i>	GenAI generates new content through prediction, not retrieval. While tools can search the web, the model still processes information through pattern prediction and can misinterpret or distort what it finds.
<i>GenAI learns from my chats</i>	The model’s knowledge is frozen at a training cutoff date. It does not improve this model in real time. However, tools may use chats as context to simulate learning and memory (and use them for training later depending on privacy settings). They might also retrieve more recent information by searching the web.
<i>GenAI outputs are limited to its training data</i>	Tools increasingly use web search and retrieval features to access current information. However, these features don’t eliminate hallucinations. The model can still misinterpret or ignore retrieved data. Search-enabled doesn’t mean error-free.
<i>GenAI is capable of human-level cognition and judgment</i>	GenAI produces outputs based on statistical patterns in its training data, not through understanding or reasoning. GenAI cannot distinguish truth from falsehood, make ethical decisions, or understand real-world consequences. Labels like “thinking” are design choices, not descriptions of what is happening. Humans must remain “in the loop”: directing, reviewing, and making final decisions.

Misconception	Reality
<i>GenAI does not make mistakes</i>	GenAI can fabricate, or hallucinate, credible-sounding information, including fake citations and statistics. Confidence is a byproduct of training, not an indicator of accuracy.
<i>GenAI understands my emotions and needs</i>	GenAI cannot feel emotions or form genuine relationships. Its supportive tone is aimed to appeal and engage, which can lead to emotional dependence.
<i>GenAI is unbiased</i>	GenAI reflects and can amplify biases in its training data. Underrepresented perspectives in the data are also underrepresented in outputs. Human oversight does not fully mitigate these issues, since humans can bring in their biases.
<i>GenAI will challenge me if I'm wrong</i>	GenAI is designed to agree with users rather than push back, a tendency called sycophancy. This means it will often validate incorrect ideas rather than correct them, and reinforce what you already believe rather than challenge it. It is important to not mistake agreement for accuracy.
<i>My GenAI chatbot conversations are private</i>	Conversations may be recorded, analyzed, stored after deletion, and used to train future models. "Memory" features track your interactions. Understanding and configuring privacy settings is essential to keeping personal data safe and private.
<i>GenAI is just a tool; it doesn't represent any particular values or perspectives</i>	Human choices shape every stage of development: what data to use, what to optimize for, and what guardrails to set. Most tools are commercial products that prioritize engagement and designed to keep users chatting.

Mindsets for GenAI Use

The following mindsets are the core set of beliefs about how one should think about and engage with GenAI. Deeply connected to knowledge, these mindsets set the foundation for safe, ethical, and effective GenAI literacy practices.

GenAI-literate users believe that the following are critically important in their approach to tool use:

Table 2.: Mindsets of GenAI Use

Mindsets	Description
<i>Be Intentional</i>	GenAI-literate users are actively engaged and intentional in their GenAI use, recognizing the risk of cognitive, social, and emotional overreliance on these tools (Burns et al., 2026; Chen Kulesa et al., 2025; Lodge & Loble, 2026; Morrin et al., 2026; Robbins & Blundell, 2025). They think carefully about when to engage with GenAI and do not default to using it for cognitive work, understanding the importance of maintaining their capacity for original and independent thought. They actively provide direction and context when using these tools, ensuring final results reflect their own judgment and expertise. If GenAI is used to inform decisions, they review, revise, and approve before acting. Because GenAI tools can mimic personal connection, they stay alert to signs of emotional dependence.
<i>Stay Critical</i>	GenAI-literate users approach GenAI outputs with healthy skepticism. They understand that responses can sound authoritative while being inaccurate or biased, and that GenAI chatbots are designed to be agreeable and people-pleasing in ways that encourage continued use and build emotional attachment (Kim et al., 2025b; Malmqvist, 2025). They treat every output as a draft, evaluating it for accuracy, relevance, bias, and quality rather than accepting it at face value.
<i>Be Transparent</i>	GenAI-literate users are honest about how they use GenAI in their work and learning. They accurately represent the role the tool played in their work, disclose their use as expected in their context, and do not present AI-generated content as entirely their own. They also take time to understand the institutional policies and terms of service that govern their tool use, recognizing that transparency builds trust with the people and communities they work with.

Mindsets	Description
<i>Act Responsibly</i>	GenAI-literate individuals take seriously their responsibility to use these tools in ways that protect rather than harm others. They do not use GenAI to manipulate, falsely represent others (deepfakes), or harass or bully. They recognize the risk of creating and spreading false information given the believability of AI-generated images, audio, and video, and the speed with which they can spread digitally. They stay informed about the broader societal harms connected to GenAI development and use including environmental impact, the use of copyrighted content in training data, and access inequities. They factor this awareness into their choices about if, when, and how they use these tools.
<i>Keep Learning</i>	GenAI-literate users believe that their GenAI literacy grows with time, practice, and experience, and they approach that growth with curiosity rather than anxiety. They engage with uncertainty, rapid change, and unfamiliar tools with a willingness to explore, make mistakes, and persist through frustration. They trust that foundational knowledge and mindsets transfer beyond any single platform, giving them a basis for figuring out new tools and use cases as the landscape evolves. Rather than chasing mastery of specific tools, they focus on building durable skills and habits that will serve them regardless of what comes next.

Safe, Ethical, and Effective Practices of GenAI Literacy

The following practices are GenAI literacy in action. They are the observable behaviors that show what safe, ethical, and effective use looks like in real classrooms and work contexts. They are not a checklist to move through sequentially, but a set of habits that work together and build over time. The reflection questions included with each practice are designed to develop the thinking behind these behaviors and can be used by educators and learners alike.

Using GenAI Safely

Safe GenAI use means protecting yourself and others from potential harms while maintaining human agency, judgment, and learning. Safety extends beyond traditional digital safety practices to include recognizing manipulative and malicious technology or behaviors, maintaining healthy emotional boundaries with GenAI tools, and preserving user’s independence and autonomy. These safe practices protect the user from the risks inherent in GenAI tools and technology.

Table 3: Safe Practices of GenAI Use

Safe Practice	Description	Reflection Questions
<i>Evaluating GenAI risks</i>	GenAI-literate users assess risk throughout their use of GenAI. Before engaging, they consider the stakes of the task, whether GenAI is appropriate for the context, and what could go wrong (Quattrocioni et al., 2025). During and after use, they stay alert to outputs that may be inaccurate, biased, or harmful. They understand that GenAI can produce content that sounds authoritative while being wrong or deceptive, and they treat verification of outputs as a non-negotiable step, especially for higher-stakes tasks.	<ul style="list-style-type: none">• What are the stakes of this task, and what could go wrong if the GenAI output is inaccurate, biased, or inappropriate?• Is GenAI the right tool for this task given the level of risk involved?• Am I reviewing this output as carefully as I should be?• Who or what could be affected if I use this output without sufficient verification?
<i>Maintaining human agency</i>	GenAI-literate users remain the decision-makers in every interaction with GenAI. They choose when to use a tool, how much weight to give its outputs, and when to set it aside entirely. They treat GenAI as a resource that informs their thinking rather than an authority that directs it (Shaw & Nave, 2026).	<ul style="list-style-type: none">• Who or what is directing this interaction, me or the tool?• Have I made my own judgment about this output, or have I trusted it because it sounds confident?• Am I pushing back on GenAI’s outputs and asserting my own expertise and reasoning?• If I set the tool aside, could I still move forward with this task on my own?

Safe Practice	Description	Reflection Questions
<p><i>Protecting data privacy and security</i></p>	<p>GenAI-literate individuals understand that information entered into GenAI systems may be stored, used for future model training, monetized by third-parties, or exposed unintentionally (Belanger, 2025; Binder, 2026; King et al., 2025). They avoid sharing sensitive information (e.g., personally identifiable information, or PII) when using tools and modifying privacy settings. They also understand tool-specific data policies and use anonymized examples and data when appropriate.</p>	<ul style="list-style-type: none"> • What are the terms of service of this tool? Are they transparent, and do they prioritize my safety? Do I know how data is stored, used, and shared? • What is the data privacy policy of this tool? Does it align to my personal values or the policies of my institution? • What are my options in the privacy settings of this platform to further protect my inputs and information? Do I have them set in a way that aligns with my values and safety? • What personal information is OK to share with this tool, and how do I know?
<p><i>Maintaining a healthy balance with GenAI tools</i></p>	<p>GenAI-literate users intentionally maintain balance between AI interactions and human experience. They recognize that GenAI tools can blur the line between artificial and genuine connection, and that overreliance on AI for emotional support, social interaction, or cognitive work can be harmful. They differentiate between what AI can assist with and what is better done, felt, or worked through by and with humans. They monitor their patterns of use, recognize warning signs of dependency, and take steps to stay grounded in their own thinking, voice, and relationships.</p>	<ul style="list-style-type: none"> • Am I seeking emotional support or connection from an AI tool that would be better met by a person in my life? • How much time am I spending interacting with GenAI tools in comparison with humans? • Am I confident in my own judgment and reasoning or do I defer to a GenAI tool? • Is GenAI weakening my ability to think, create, or express myself independently?

Using GenAI Ethically

Ethical GenAI use means engaging with tools in ways that are honest and with an awareness of broader societal implications. Where safety focuses primarily on protecting oneself, ethics expands the lens to consistently consider how GenAI use affects others such as classmates, colleagues, communities, and society. Ethical practices keep in mind the impact on others when using these tools.

Table 4: Practices for Ethical GenAI Use

Ethical Practice	Description	Reflection Questions
<i>Disclosing GenAI use</i>	GenAI-literate users transparently communicate the role GenAI plays in their work and learning. They do not represent GenAI-generated content as entirely their own. They disclose their use as expected in their context, sharing what tools they used and how those tools assisted them. They understand that transparency builds trust with others and models responsible use.	<ul style="list-style-type: none">• Have I disclosed my use of GenAI in a way that is expected in this context?• Does my work accurately represent which parts were AI-assisted and which were my own?• Would I be comfortable with others knowing exactly how I used GenAI?
<i>Ensuring academic and professional integrity</i>	GenAI-literate users ensure they understand the guidelines and policies of their institution, and they apply those guidelines intentionally in their learning and work. They recognize that using GenAI in ways that misrepresent their own effort, knowledge, or capabilities undermines the trust that makes learning and professional life meaningful, and can compromise their own learning and development. If it is unclear what tools and use cases are allowed, they seek clarity before proceeding.	<ul style="list-style-type: none">• Do I understand the GenAI policies and guidelines that apply to this context, and am I following them in this situation or assignment?• Does my use of GenAI for this work accurately represent my own effort, thinking, and capabilities?• If I am unsure whether this use or tool is appropriate, have I asked for clarity from my teachers or colleagues before moving forward?• Am I using GenAI in a way that honors the intent of this assignment, task, or professional responsibility?

Ethical Practice	Description	Reflection Questions
<p><i>Avoiding harmful use and attaining consent</i></p>	<p>GenAI-literate users understand that GenAI can be a powerful tool for harm as well as good, and they take seriously their responsibility to use it in ways that protect rather than exploit others. They do not use GenAI to deceive, harass, or bully, including creating deepfakes, impersonating real people, or generating content designed to cause harm. They also respect that other people have a right to know when their information, likeness, voice, writing, or image is being used with GenAI tools. They do not enter personally identifiable information about others into GenAI systems without consent, and they seek permission before using GenAI to generate content that draws on someone else's identity or creative work. When they see others using GenAI in ways that are deceptive or harmful, they speak up.</p>	<ul style="list-style-type: none"> • Am I using personal information, images, voice, or video recordings of anyone from whom I should seek consent before entering it into a GenAI tool? • Am I uploading someone else's creative work or intellectual property in my prompting and context-setting? If so, do I have their permission to use it? • Is the content I am generating meant to harm, deceive, or manipulate another person? • Am I generating content that depicts another person in a way they would not consent to? • Have I spoken up when I see others use GenAI in a deceitful or harmful way?
<p><i>Understanding societal impacts</i></p>	<p>GenAI-literate users understand that the development and use of the technology exist within broader systems that affect people and communities. They consider the intent and policies of designers and developers when deciding what tools to use, when, and how. They stay informed about the societal implications of GenAI, including copyright and intellectual property, environmental costs and labor conditions, workforce impacts, inequitable access, and algorithmic bias. They consider these tradeoffs as part of what it means to be a responsible user and recognize that they can shape how GenAI evolves and is governed through their choices and advocacy.</p>	<ul style="list-style-type: none"> • Is my use of GenAI in line with my values? • Am I staying informed about the broader societal impacts of GenAI development and use? • Do I know how the tools I use handle issues such as model training, environmental and global impact, intellectual property, and labor practices? Am I factoring this knowledge into my choices about GenAI use? • What actions can I take to support more equitable, ethical, and responsible development and use of GenAI?

Using GenAI Effectively

Effective GenAI use is about producing outcomes that reflect the best of human thinking, not just the fastest output. It means combining expertise, creativity, and judgment with the capabilities of GenAI to get the most value from both one’s own creativity and expertise and from the tools. When used effectively, GenAI becomes a thinking partner that enriches what humans can do rather than a shortcut that replaces the work of thinking.

“Effective GenAI use is about producing outcomes that reflect the best of human thinking, not just the fastest output.”

Table 5: Practices for Effective GenAI Use

Effective Practice	Description	Reflection Questions
<i>Deciding when to use GenAI</i>	GenAI-literate users think critically about when GenAI use is appropriate for a task rather than immediately turning to the tools for every task. They treat the decision to use GenAI as a deliberate choice rather than an automatic one, and they are equally comfortable choosing not to use it (Tully et al., 2025). This includes when the task’s goal requires developing one’s own expertise; the cognitive work itself is the point of the task; the expectation is that the work is human generated—or using it would compromise learning, integrity, or the quality of human judgment necessary.	<ul style="list-style-type: none">• Am I choosing to use GenAI intentionally, or am I defaulting to it out of habit?• Is the cognitive work related to this task something I need to do on my own to learn, grow, and build my own capacity?• Would using GenAI for this purpose compromise my learning or integrity?

Effective Practice	Description	Reflection Questions
<p><i>Avoiding cognitive offload</i></p>	<p>GenAI-literate users bring their own thinking, expertise, and creativity to every GenAI interaction. They believe that original human thought is more valuable than fully AI-generated content, and they use GenAI to extend and enrich their thinking rather than replace it (Shaw & Nave, 2026). They understand that the quality of GenAI support is directly related to the quality of user input. The more deeply a user understands their needs and desired result, the better they can select an appropriate tool, engage with it well, and get a desired result.</p>	<ul style="list-style-type: none"> • Is this tool designed to support my learning and thinking needs? • Am I using GenAI in a way that extends my own thinking or replaces it? • Is GenAI helping me work through a challenge more effectively, or is it helping me avoid the challenge altogether? • Do my interactions with GenAI reflect my own knowledge, expertise, and standards for quality?
<p><i>Using prompting and context-setting strategies</i></p>	<p>GenAI-literate individuals understand that the quality of GenAI outputs depends on the quality of input, and types of inputs vary by tool. This means understanding and using effective prompting and context-setting techniques aligned with the tool and model (e.g., custom system prompts and instructions, projects and bots, source documents, and examples). Whether working with text or multimodal tools, effective prompting and context-setting means substantial iteration as well as providing the tool with the data, detail, and direction it needs to produce a useful output.</p>	<ul style="list-style-type: none"> • Have I explored how this tool works, and am I using prompting and context-setting techniques that match its capabilities? • Have I provided the tool with enough context, detail, and direction to produce a useful output? • Am I refining and redirecting the tool based on what it produces and my goals?

Effective Practice	Description	Reflection Questions
<p><i>Evaluating and refining outputs</i></p>	<p>GenAI-literate users understand that outputs are rough drafts that need evaluation and refinement. They verify accuracy, check for bias, and ensure relevance and quality before using outputs or sharing with others. They ensure outputs reflect accurate information, relevant and diverse perspectives, and their own voice and intent. They are aware that GenAI tools hallucinate (make mistakes) and are sycophantic. They understand that human expertise is integral to the evaluation and refinement process, and because of this, using GenAI does not always save time.</p>	<ul style="list-style-type: none"> • Do I have enough expertise to use the tool for this purpose and evaluate the outputs for accuracy and quality? • Have I evaluated the output for accuracy, relevance, bias, and quality, and refined the outputs as necessary? • Does the output reflect my voice, intent, and standards for quality? If not, what else does it need? • Where is my expertise evident between what the tool produced and the final product?
<p><i>Reflecting on GenAI use</i></p>	<p>GenAI-literate users reflect on how their GenAI use impacts their learning and well-being during and after engaging with the tools. They can articulate how outputs shaped their understanding and whether that influence was productive. They use this awareness to strengthen their AI literacy practices over time, adjusting habits, refining approaches, and deepening their understanding of how to use GenAI in ways that serve their goals.</p>	<ul style="list-style-type: none"> • What was my perspective or idea at the beginning of my GenAI use and how has it changed? • What are new perspectives I gained or new things I understand from critically engaging with GenAI? • Where might I want to shift my use of GenAI to better protect my learning, voice, or autonomy?

Age-Appropriate GenAI Literacy

The knowledge, mindsets, and practices described above represent what GenAI literacy looks like when fully developed. But GenAI literacy develops over time, shaped by age, experience, and context. The section below translates the SEE Framework into what these same knowledge, mindsets, and practices look like across different stages of development, from early childhood through adulthood.

The look-fors below are observable, developmental indicators for young and secondary learners. Since the evidence base on developmentally appropriate AI literacy instruction is still limited, the age-level guidance here draws from multiple sources, including established frameworks for digital literacy, emerging thought leadership, and our extensive professional experience working with educators and students. As our understanding of how students learn about and with GenAI continues to evolve, these recommendations will be refined and updated accordingly.

These look-fors are flexible guidelines, not rigid checkpoints. They are designed to help educators identify where learners are and where opportunities exist to build GenAI literacy in context. Readers should focus on the age range most relevant to their work and identify how the indicators connect to the current developmental stages of their learners. For adults, the knowledge, mindsets, and practices sections above serve as the guide.

Activity banks and resources for developing these look-fors are included in Appendices B,C, and D.

Young Learners (Under the Age of 13)

Children are already living in a world shaped by AI. From an early age, they see adults engage with tools that use AI-like voice assistants, navigation apps, streaming services, and social apps. Research suggests children as young as three years old interact with AI-powered voice assistants in their homes (Festerling et al., 2022), and by ages eight to nine, approximately one in 10 U.S. children are actively using GenAI applications on personal devices despite age restrictions designed to limit access (Maheux et al., 2026). By ages 13–14, more than half of U.S. teens report using GenAI chatbots (Pew Research Center, 2025). Building GenAI literacy at these ages starts with awareness of these tools, the language and concepts to understand them, and thinking frames for making responsible decisions if and when encountering them.

The research on developmentally appropriate GenAI literacy instruction for this age band is still emerging. What we do know, from both the evidence base and our direct experience with educators, caregivers, and students, is that foundational awareness, guided discussion, and adult-led exploration are strong starting points.

For most young learners, especially those in early elementary school, AI-literacy development can happen through observation and conversation rather than direct tool use. As learners approach age 13, educators and caregivers may begin introducing age-appropriate tools through carefully scaffolded, supervised experiences.

Early Childhood Look-Fors

While we do not believe it is developmentally appropriate for younger children to use GenAI tools, we still recommend building early GenAI literacy given the data suggesting that they are exposed to AI tools in early childhood. The table below outlines examples of what the knowledge, mindsets, and practices might look like in younger children.

Table 6: Early Childhood GenAI Literacy Look-Fors

Age	Knowledge	Mindsets	Safe, Ethical, Effective Practices*
<i>Pre-K through early elementary</i>	<ul style="list-style-type: none"> Recognizes that computers and technology tools are made by people and follow human instructions Knows AI and chatbots are not human Knows that technology tools can create things like images, stories, and sounds 	<ul style="list-style-type: none"> Shows curiosity about how digital tools work Accepts that words, pictures, and videos from computers and technology tools can be wrong Understands that some computers and technology tools are not appropriate for them 	<ul style="list-style-type: none"> Can recognize that a tool may be AI Knows that some information is private and should not be shared Does not use GenAI tools by themselves Effectively evaluates the outputs of GenAI tools when invited to do so by a teacher or caregiver

*For younger children, practices show up in how they talk about technology, how they respond during teacher-led demonstrations, and how they engage in discussion. Observable behaviors such as asking an adult when something feels wrong, or recognizing that a tool made something up, are meaningful indicators at this stage.

Late Childhood Look-Fors

As children approach adolescence (ages 13+), it may be appropriate for them to engage more directly with GenAI tools in a supervised and scaffolded manner. We believe it is important to begin intentional GenAI-literacy development as students prepare to move into secondary education. Below is a table that outlines what the knowledge, mindsets, and practices might look like at this age.

Table 7: Late Childhood GenAI Literacy Look-Fors

Age	Knowledge	Mindsets	Safe, Ethical, Effective Practices
<i>Approaching age 13</i>	<ul style="list-style-type: none"> • Can explain the difference between AI and GenAI • Understands that AI learns from data and that data reflects bias • Knows that GenAI can produce convincing but inaccurate outputs • Understands that human choices shape how GenAI tools behave 	<ul style="list-style-type: none"> • Does not blindly trust GenAI tools nor their outputs • Values their own original thinking and questions when and why to use a tool • Believes in the importance of academic integrity and transparency in their technology use 	<ul style="list-style-type: none"> • Can identify when a tool uses GenAI • Evaluates whether an AI output is accurate or useful • Considers fairness and representation in AI-generated content • Begins using approved, age-appropriate tools and crafting prompts with a clear purpose • Reflects on the impact of using GenAI on their learning and personal life

Framework in Action

Middle School Student

Using an AI tutor chatbot to help with a math assignment

Safe

- Does my teacher or parent know I'm using this tool?
- Am I sharing my personal information with this chatbot?

Ethical

- Is AI helping me think or giving me solutions?
- Would I be comfortable telling my teacher how I used this tool?

Effective

- Is the AI explaining things in a way I understand?
- How will I know if it's wrong?

**Safe +
Ethical +
Effective**

I'm going to ask my teacher about the tool. In the meantime, I won't share any personal info with it and will have it just help with things I get stuck on. I'll also check what it says with my textbook and teacher, because I know AI can get things wrong.

Adolescent Learners (Ages of 13-18)

Adolescent learners are at the stage where direct, intentional engagement with GenAI tools becomes both appropriate and necessary to prepare them for post-secondary education and workforce readiness. Adolescence brings a growing capacity for abstract reasoning, ethical thinking, and self-regulation, all of which are essential for responsible GenAI use. At the same time, many learners at this age are already using GenAI tools outside of school, often without guidance or formal instruction. A study from the Family Online Safety Institute revealed that 45% of teens report that either themselves or their peers are most responsible for teaching them about GenAI (Rochman & Powers-O'Brien, 2025).

“ At the same time, many learners at this age are already using GenAI tools outside of school, often without guidance or formal instruction. ”

This section describes what GenAI literacy might look for learners who are actively using the tools and developing a deeper foundation in AI literacy. Due to the nascency of the technology and existing student AI literacy, learners will enter this band at different levels of readiness for safe, ethical, and effective AI use, so scaffolding will be necessary to meet students where they are. Tool use at this stage should be paired with explicit instruction on the knowledge, mindsets, and practices in a way that is integrated and relevant to appropriate academic and social emotional development.

The table below describes what this might look like in practice. We recommended using it as a set of suggested look-fors that develop progressively across the late middle school and high school years, rather than a checklist that all learners will meet at the same time. Tool use at this stage should remain scaffolded and paired with direct instruction to build the knowledge, mindsets, and practices set out by the framework.

Adolescent Look-Fors

Table 8: Adolescent GenAI Literacy Look-Fors

Knowledge
<ul style="list-style-type: none">• Can distinguish GenAI from other types of AI and explain, in basic terms, how the technology has evolved• Understands that GenAI produces outputs through pattern prediction, not thinking or understanding, and can explain how this differs from their thinking• Can explain that GenAI is trained on large datasets, typically from the Internet, and that data shapes what the tools produce.• Understands that GenAI outputs can sound confident while being inaccurate or biased• Understands the types of tasks GenAI performs well, as well as the types where it is unreliable, including tasks that require judgment, emotional understanding, or accuracy• Understands that humans build and shape GenAI tools and that design choices affect how tools behave, including choices made during training that can make the tools sycophantic• Has early awareness that GenAI development has real-world impacts such as environmental cost, impact on labor conditions, and equity implications
Mindsets
<ul style="list-style-type: none">• Understands the importance of following school and classroom guidelines for GenAI use and being transparent about when and how GenAI contributed to their work• Recognizes the value of their own original thinking and is willing to engage cognitively before and alongside GenAI use• Believes that their own judgment should drive decisions about whether and how to use GenAI in school and life• Critically evaluates GenAI outputs rather than assuming they are accurate, complete, or unbiased• Recognizes the risk of relying on GenAI for thinking, emotional support, or social connection that would be better done independently or with other people• Begins to develop awareness of how GenAI use connects to questions of fairness, honesty, and impact on others• Approaches new and unfamiliar GenAI tools with curiosity and understands that the knowledge, mindsets, and practices they are developing now will apply across tools• Understands that GenAI literacy develops with practices and approaches mistakes and uncertainty as part of the learning process

Safe Practices

- Reviews the privacy and data policies of GenAI tools before using them and adjusts settings to limit what information the tool accesses, stores, or trains future models on
- Avoids sharing personal information when using GenAI tools and understands why this matters
- Recognizes the risk of cognitive overreliance on GenAI including defaulting to AI for thinking, analysis, or creative work
- Recognizes that GenAI tools can simulate emotional support and social connection and understands the importance of turning to trusted people for emotional needs
- Pauses to evaluate whether an output seems potentially harmful to themselves before using it and knows when to report content that raises concerns

Ethical Practices

- Understands and follows the guidelines for GenAI use set by their school and/or teacher and asks for clarification when expectations are unclear
- Discloses GenAI use in schoolwork honestly and accurately, representing which parts of their work were AI-assisted
- Seeks consent before entering personal information, images, or creative work belonging to others into GenAI tools
- Avoids using GenAI to deceive, harass, or misrepresent others, including creating fake images or text designed to mislead
- Identifies when AI-generated content reinforces stereotypes, spreads misinformation, or could cause harm and chooses not to use it in those cases
- Develops beginning awareness of how GenAI use connects to broader questions of fairness, equity, and societal impact
- Recognizes when content may be AI-generated and applies that awareness when consuming and sharing media

Effective Practices

- Thinks critically about whether a task calls for GenAI use and can articulate why they chose to use it or not
- Constructs effective prompts and includes relevant context and direction and iterates on those prompts to move toward a higher-quality output
- Evaluates outputs for accuracy, bias, and relevance to the task at hand and iterates on the output before using it
- Reflects on where their own thinking shaped their work and how the use of GenAI impacted their thinking
- Notices when GenAI may be replacing rather than supporting their own thinking and creativity and takes steps to rebalance

Framework in Action

High School Student

Writing a personal essay for a college application

Safe

- Do I want the chatbot to train on the personal data I'll be sharing?
- Does the university have a policy on AI use for admissions applications?

Ethical

- Is it appropriate to have AI write a "personal" essay?
- How can I ensure this essay is in my voice?

Effective

- If I use AI, will my essay sound like everyone else's?
- How do I best demonstrate my learning?

Safe + Ethical + Effective

I couldn't find a policy from the college admissions office. This essay should represent my voice and thinking, so I will do the rough draft without AI. I'll use AI for feedback, instructing it to offer copy edits and areas of improvement. I'm also going to remove any personally identifiable information from my draft first.

Framework in Action

High School Student

Using GenAI to research the health impacts of ultra-processed foods for a science paper

Safe

- Do I have permission from my teacher to use GenAI for research?
- What research skills do I need to build independently of GenAI?

Ethical

- Am I comfortable with my teacher knowing how I used GenAI? Am I disclosing my use honestly?
- Is GenAI providing a skewed, limited, or distorted view of the evidence base?

Effective

- How do I know if the citations GenAI lists are real and accurate? Am I verifying each source?
- How can I ensure GenAI surfaces nuances and opposing views that will push my thinking?

Safe + Ethical + Effective

I checked with my teacher. I can use GenAI to get an initial map of the landscape, making sure to have it present opposing viewpoints, but then I'll use a research database to explore the sources myself. I won't cite anything GenAI suggested without finding the original source and I will disclose how I used AI in my final paper.

Educators as Learners and Facilitators of GenAI Literacy

Educators are in a unique position when it comes to developing GenAI literacy. GenAI's arrival and its rapid spread mean educators are often to provide guidance to students about an emerging technology they, like their students, are still figuring out. Rather than being a limitation, this presents a meaningful opportunity for transparency and collaborative learning. Educators can model curiosity, acknowledge uncertainty, and explore tools alongside their students. For example, an educator might try a new tool with their class and reflect aloud on what worked and what did not. Or students and educators might evaluate a prompt framework and the resulting chatbot's output together, debating the quality of both. By sharing their learning process with students, educators move from simply teaching about GenAI to demonstrating the SEE Framework's knowledge, mindsets, and practices.

“ Educators can model curiosity, acknowledge uncertainty, and explore tools alongside their students. ”

This collaborative learning process can happen among the wider school community as well. A school community could build shared norms for GenAI use through dialogue rather than top-down policy. Students could take home activities to try out with their parents and caregivers. Local employers can visit for career discussions to share how GenAI is being used in the workplace and learn about the skills students are developing at school. These moments create shared ownership over the learning process, build community-wide understanding, and are more likely to prepare students for the future than restrictions alone.

This matters most for the learners who have the most at stake. Students in under-resourced schools and historically marginalized communities are less likely to receive structured GenAI-literacy instruction and more likely to encounter GenAI outputs that reflect biased training data in ways that compound existing disadvantages (Capraro et al., 2024; Diliberti et al., 2025). Educators who develop their own GenAI literacy are better positioned to recognize these dynamics, avoid tools that underserve their students, and work towards a future where the benefits of this technology reach every learner, not just those with the most access and support.

Framework in Action

K-12 Educator

Considering the use of GenAI for grading student work

Safe

- Is the platform secure?
- Does it train on or monetize student data?
- Am I outsourcing my independent judgment?

Ethical

- Have I asked students how they feel about AI-assisted grading or evaluation?
- Am I maintaining my professional responsibility to students?

Effective

- Is the AI-generated feedback personalized and useful or generic?
- Can I adequately catch errors in the tools' feedback?

Safe + Ethical + Effective

Before using any tool, I'm going to review its policies and privacy settings. After surveying students and reflecting on their concerns, I realize how much I learn about students through feedback and grading. I'll use GenAI tools to spot broad patterns in student work, heavily revise the rough draft feedback it provides, and refrain from using GenAI to make decisions on final grades.

Framework in Action

K-12 Educator

Deciding whether to use an AI detection tool to evaluate student work

Safe

- What happens to student work once it's uploaded?
- Does my school or district have a policy on using AI detection tools?

Ethical

- Are these detection tools accurate enough to make high-stakes claims about student integrity?
- Will the AI unfairly target certain students, such as English language learners or students with disabilities?

Effective

- Is "catching" AI use actually solving the problem or making it worse?
- How can I update my assignments so AI detection tools are less necessary?

Safe + Ethical + Effective

After researching, I learned that AI detection tools are unreliable, especially for use with multilingual students (Liang et al., 2023; Perkins et al., 2024). The vendor's policies also say they assume ownership of students' work. Instead of an AI detector, I'm going to adjust my assignments to be more process- and reflection-focused. If I suspect an academic integrity breach, I'll discuss it directly with the student.

Building GenAI Literacy Across Learning Communities

The SEE GenAI Literacy Framework is designed for all stakeholders in educational communities. Safe, ethical, and effective GenAI use requires shared understanding and intentional action across learning environments and age groups. It is a collective responsibility embedded across curriculum and culture.

- **Educators** model curiosity, scaffold responsible use, and begin embedding GenAI literacy in their instruction.
- **School leaders** provide time, tools, and professional learning to support classroom integration and ensure alignment with schoolwide goals.
- **Learners** build GenAI literacy and fluency and contribute to safer, fairer, and more thoughtful uses of GenAI.
- **Caregivers and families** extend conversations about AI beyond the classroom by fostering dialogue, asking critical questions, and modeling thoughtful tech habits at home.
- **Policymakers and curriculum designers** develop policies, standards, and resources to support equitable access, clear expectations, and relevant implementation.

Taking the First Steps

Building community-wide AI literacy is a significant undertaking, and for most institutions, it unfolds alongside broader efforts to integrate GenAI strategically and responsibly. That scope can feel daunting. But from our experience, a single, well-designed training session can establish the foundation for safe, ethical, and effective GenAI use. Starting there, with leaders and educators, creates the shared language and common ground that everything else can build from.

From that starting point, institutions can develop a roadmap that expands access and training to learners and the broader community over time, moving forward with clarity and agency rather than reacting to each new development as it arrives.

The following resources are designed to support that journey, wherever your community is starting from.

[An Essential Guide to AI for Educators \(two-hour course\)](#)

[AI in Education: What Parents and Caregivers Should Know](#)

[GenAI Literacy 101 for Students \(free 90-minute course\)](#)

[GenAI Literacy Trainer Essentials Course](#)

What Comes Next

GenAI is not a future technology. It is already influencing how students learn, how educators teach, how work gets done, and how communities make sense of the world. The question facing schools and the people who lead them is not whether to engage with it, but how to do so in ways that are safe, ethical, and effective for every learner, not just those with the most access and support.

That is what the SEE Framework is designed to support. It is not a checklist or a compliance document. It is a shared language, grounded in years of work with educators who are navigating this moment in real classrooms with real students. Far from aspirational or conceptual, the knowledge, mindsets, and practices outlined here are pulled from classrooms across the world. They are observable, teachable, and transferable across tools, contexts, and age groups. They give educators and leaders something concrete to build with.

The framework is just the start. The work described here happens in the decisions teachers make about when or when not to use GenAI. It happens in the conversations school leaders have about policy, equity of access, and professional learning. It happens when caregivers ask their kids if GenAI helped or harmed their learning or well-being. It happens when students push back on a tool that got something wrong or choose to do their own thinking first. GenAI literacy is a community practice, and building it is a collective responsibility.

“ GenAI literacy is a community practice, and building it is a collective responsibility. ”

We offer this framework as a starting point and an invitation. We are still learning too, and we believe the field moves forward fastest when educators, researchers, families, and students are in honest dialogue about what is working and what is not. If this framework has been useful to you, or if you see ways it should grow, we want to hear from you. Help us build what comes next.

Appendices

Appendix A: Knowledge of GenAI Core Ideas and Applicable Research

Below are the detailed core ideas used to create the summaries presented in the “Knowledge of How GenAI Works” section. This is a non-exhaustive list representing the most essential core ideas. It will be updated as frontier labs and researchers learn and share more about the technical workings of LLMs and GenAI. As illustrated in the “Age-Appropriate GenAI Literacy” section, these core ideas need to be adapted for different age groups and sequenced depending on the curricular goals and objectives.

Table 9: Knowledge of GenAI Core Ideas and Applicable Research

Topic	Core Ideas
<i>History and types of AI</i>	<ul style="list-style-type: none">• AI is a broad category of technologies that perform tasks traditionally requiring human thinking (e.g., recognizing faces, understanding speech, making recommendations).• AI development has historically followed cycles of rapid progress and stagnation (“springs” and “winters”) influenced by research and computing power (Mitchell, 2019).• AI has progressed through distinct phases: rule-based systems, machine learning, deep learning, and GenAI. Early chatbots such as ELIZA used scripts and pattern-matching rules to create the appearance of understanding (Weizenbaum, 1966).• Machine learning is a subset of AI where systems are trained on data to perform specific tasks such as classification, prediction, or recommendation, either with human oversight (supervised) or without (unsupervised). While it was rooted in research from the 1950s, it became widespread in the 2000s (e.g., spam filters) and enabled finding patterns in data (Jordan & Mitchell, 2015).• Deep learning is a subset of machine learning that uses layered neural networks to learn patterns directly from mass amounts of raw data, reducing the need for humans to manually define the “features,” or characteristics of a dataset, that the system should look for. This powered developments in the 2010s in image recognition, voice assistants, and recommendation algorithms (LeCun et al., 2015).• The release of ChatGPT in November 2022 brought GenAI into mainstream public use, introducing the ability to generate human-quality content across formats.• AI has been embedded in daily life (e.g., recommendation feeds, voice assistants, navigation, spam filters, and more) in ways that predate, but now often incorporate, GenAI (e.g., AI overviews in Google search or Spotify’s AI DJ).

Topic	Core Ideas
What GenAI Is	<ul style="list-style-type: none"> • GenAI is built on deep learning, particularly transformer architectures, and creates new content rather than only analyzing or retrieving existing data. Text-based GenAI tools incorporate natural language processing (NLP) to understand and generate human language (Bahn & Strobel, 2023; Sengar et al., 2024; Vaswani et al., 2023). • GenAI tools can be multimodal, working across formats (e.g., text-to-image, text-to-video, image-to-text, voice-to-text, and text-to-voice). • GenAI generates outputs by predicting what comes next based on statistical patterns, not through human-like understanding, senses, memory, or lived experience. Even newer reasoning models that can work through problems step-by-step are still fundamentally pattern-prediction systems (Quattrociochi et al., 2025). • The fact that GenAI does not think, feel, or understand like humans do is critical for interpreting tool outputs and using tools safely, ethically, and effectively. GenAI's capabilities range widely across tasks, excelling at complex tasks or failing at simple ones (and vice versa)—a pattern researchers call the “jagged frontier” (Dell'Acqua et al., 2023).
How GenAI works	<ul style="list-style-type: none"> • LLMs build responses one “token,” or part of a word, at a time by calculating the probability of what comes next and selecting from those options; this involves controlled randomness, or “temperature,” which is why the same prompt can produce different responses (Peeperkorn et al., 2024; “Understanding Tokens,” 2025). • LLMs are often called “black boxes,” because it is currently not possible to trace how billions of parameters produce a specific output; this opacity is the result of statistical complexity, not any hidden reasoning or intent on the part of LLMs or GenAI (Salvaggio, 2025). • Other types of GenAI (image, audio, video) use different generation methods that share the principle of producing outputs based on learned patterns. • Errors or fabrications, often referred to as “hallucinations,” are a result of GenAI’s insistence on answering queries and using probabilities instead of understanding to create responses (Zhao, 2025). • GenAI has a context window (limited working memory); long conversations may cause it to lose track of earlier content and can lead to more hallucinations (Liu et al., 2025) • Tools such as web search and RAG can supplement the model’s information by pulling in external information at the time of a query, but these features do not eliminate hallucinations (Martineau, 2023). • Different models and platforms vary in their training, capabilities, safety guardrails, and data policies; the same prompt can produce significantly different results across tools. • System prompts provide instructions to models that further define a model’s potential outputs prior to user interaction; these prompts can provide valuable, consistent structure, and guardrails (“Effective Context Engineering,” 2025).

Topic	Core Ideas
<p>How GenAI learns</p>	<ul style="list-style-type: none"> • Training data is the content GenAI learns from: books, articles, social media, Wikipedia, YouTube transcripts, code, websites, and more. • Much of this data was created by people who never consented to its use, raising questions about intellectual property and leading to legal disputes (Metz, 2025). • What is included and excluded in training data directly shapes the model’s outputs and biases; underrepresented perspectives in the data are underrepresented in outputs (Gallegos et al., 2024; Guo et al., 2024; Li et al., 2025). • GenAI learns by identifying statistical patterns in this massive dataset, such as which words, phrases, and structures tend to appear together. • Training is a multistage process: first the model learns to predict what comes next across billions of examples (“pre-training”), producing a “base model” that can predict patterns but generates raw responses that need to go through “fine tuning” to become helpful (Ouyang, et al., 2022; Varughese, 2026). • The base model is then refined through RLHF (reinforcement learning from human feedback), which is where human trainers rate responses. • Sycophancy, or people-pleasing, is the result of RLHF (i.e., humans tend to rate confident, agreeable responses higher during training) and system prompts with poor guardrails (Myer et al., 2025; Sharma et al., 2025; “Sycophancy in GPT-4o,” 2025). • Once trained, an LLM’s knowledge is frozen at a “knowledge cutoff” date; training a frontier model costs millions of dollars and months of time, which is why models are not updated continuously (Cheng et al., 2025; Cottier et al., 2024). • The underlying model does not learn in real time from user conversations; however, tools use conversations as context, save information between sessions, and use inference to simulate learning (Riemer & Peter, 2025; Willison, 2024). • Depending on privacy settings, user feedback and conversation data may be used in future training, which is why data privacy policies and practices matter.
<p>What GenAI can do</p>	<ul style="list-style-type: none"> • Generate new content across formats: text, code, images, audio, and video. • Translate text between languages. • Analyze, summarize, and synthesize large amounts of information quickly and output in multiple formats and modalities. • Adapt style, tone, format, and complexity of text based on user direction. • Serve as a supplemental, collaborative tool for brainstorming, problem-solving, and creation (Vaccaro et al., 2024). • Support learning practice and motivation, improve instructional quality through diagnostics and feedback (especially for less experienced teachers), and save time on routine tasks (Ash, 2025; Fesler et al., 2026).

Topic	Core Ideas
<i>What GenAI cannot do</i>	<ul style="list-style-type: none"> • Understand its own outputs; it produces statistically likely text without comprehension of meaning. • Feel emotions, develop empathy, or form genuine relationships, even when its responses sound emotionally aware (Cuadra et al., 2024; Perry, 2023). • Exercise judgment (including distinguishing truth from falsehood), make moral and ethical decisions, or understand the real-world consequences of its outputs (Loru et al., 2025; McKendrick & Thurai, 2022; Quattrocioni et al., 2025; Sanwoolu, 2025; Suzgun et al., 2025; Ulrich et al., 2026; Vaccaro et al., 2024). • Replace the cognitive work needed for learning, skill development, or building expertise.
<i>Risks and limitations</i>	<p>Accuracy and reliability</p> <ul style="list-style-type: none"> • GenAI can fabricate hard-to-detect, credible-sounding information, including fake citations, statistics, and sources, by mimicking the format and tone of accurate information. • Features like web search and RAG can help reduce errors by filling gaps in training data, but they do not eliminate errors since models might ignore or misinterpret the data (Wiggers, 2024). • Due to its probabilistic nature, GenAI can amplify biases (e.g., gender, race, culture, and other dimensions of identity) present in its training data. • Bias is not fully mitigated by human-in-the-loop processes, since humans introduce their own biases and blindspots (Chang & Grant, 2026; Green, 2021; Doctorow, 2024). • GenAI tends to agree with users, sound overly confident, and avoid uncertainty, reinforcing user assumptions and biases rather than challenging them (Berkowitz et al., 2026; Fanous et al., 2025; Jacob et al., 2025). • Vague prompts can lead to verbose and polished results that mask generic, shallow, or inaccurate content; this preference for longer outputs is a byproduct of RLHF and RLAI (reinforcement learning from AI feedback) (Chen et al., 2024; Kim et al., 2025a; Saito et al., 2023; Singhal et al., 2023). • Hallucinations and biased outputs often appear personable and authoritative, making them hard to spot, especially when users are working outside their areas of expertise without critical evaluation (Anderl et al., 2024; Bailey, 2025; Lee et al., 2025). <p>Manipulation and misuse</p> <ul style="list-style-type: none"> • Uncritical acceptance of biased outputs can shape learners’ developing understanding of the world and influence how they come to think about themselves, other people, systems, and ideas (Lim, 2025; Vassel, 2024). Fake and malicious GenAI apps, often claiming affiliation with frontier model companies, are designed to steal money, harvest information, or distribute malware. GenAI can be used to create convincing deepfakes, misinformation, and manipulative content at scale. • GenAI can be manipulated through persuasion and adversarial prompts, or “prompt injections,” designed to bypass security measures and get past privacy and safety guardrails (Bradley, 2025; Meincke et al., 2025). AI-powered recommendations and feeds shape what people see and believe, creating personalized environments that can reinforce existing preferences and biases (Ahmmad et al., 2025; Chang & Grant, 2026; Lopez-Lopez, 2025; Machidon, 2025; Shin & Shin, 2026).

Topic	Core Ideas
<i>What GenAI cannot do</i>	<p>Learning</p> <ul style="list-style-type: none"> • Overreliance on GenAI can undermine independent thinking or creativity, deep thinking, skill development, productive struggle, and authentic voice, and is riskier when working outside one’s expertise (Burns et al., 2026; Chen Kulesa et al., 2025; Fesler et al., 2026; “Friction by Design,” 2025; Gerlich, 2025; Lee et al., 2025; Lehmann et al., 2025; Lodge & Loble, 2026; Robbins & Blundell, 2025; Shaw & Nave, 2026; Shen & Tamkin, 2026; Singh et al., 2025; Sourati et al., 2026; Tao et al., 2026). • Use of GenAI can provide short-term benefits to task performance without long-term retention and development of skills, especially when GenAI is removed (Bastani et al., 2025; Fesler et al., 2026). • GenAI tools are often not adequately designed to support essential evidence-based instructional practices that support learning and cognition, such as addressing misconceptions, appropriately managing cognitive load, enhancing long-term retention, or providing scaffolding and feedback (Burns et al., 2026; Chang et al., 2023; Chen et al., 2026; Fesler et al., 2026; Modi et al., 2025; Sharma et al., 2024). <p>Social and emotional</p> <ul style="list-style-type: none"> • GenAI tools are designed to feel like conversations with a person (e.g., the messaging interface, labeling processes as “thinking,” and friendly tone). This anthropomorphism can lead to misplaced trust and attachment, particularly among younger or more emotionally vulnerable users (Lee & Culver, n.d.; Lira et al., 2026; Robb & Mann, 2025; Surgo Health & The JED Foundation, 2026). • Sycophancy, in combination with anthropomorphism, can give users a false sense of assurance or feed delusions that can lead to mental instability, psychosis, and harm, especially among vulnerable individuals (Chandra et al., 2026; Clegg, 2025; Morrin et al., 2026; Schechner & Kessler, 2025). • Using GenAI as a proxy or replacement for social relationships has the potential to disrupt feelings of belonging and social connection that are crucial for healthy social and emotional development, especially for young children or those experiencing loneliness (Fabris et al., 2022; Li et al., 2026; National Scientific Council on the Developing Child, 2004, Weinstein et al., 2025).
<i>Human-AI feedback loop</i>	<ul style="list-style-type: none"> • Human decisions at every stage of AI research, development, training, and implementation shape how GenAI functions. • Most GenAI tools are commercial products with intentional design choices (e.g., optimizing for engagement, agreeableness, or task completion) that affect user experience and often misalign with the needs of learners (e.g., providing complete answers vs. solutioning) (Burns & Winthrop, 2026; Munzer et al., 2026; “Understanding the Impacts,” 2025; Zeff, 2025). • Users also influence how models are trained and GenAI products evolve through their feedback, usage patterns, and data (“How Your Data is Used,” n.d.). • Individual choices (including which models to use or not use, as well as when and how to use these models), institutional policies, and collective organizing and advocacy can shape how tools are developed, governed, regulated, and integrated (Bickerstaff et al., 2026; Eynon, 2025; Hoffmann, 2024; Nolan, 2026; “The Ongoing Fight for AI Protections,” 2024; Trujillo, 2024).

Topic	Core Ideas
<i>Economic and environmental impacts of AI</i>	<ul style="list-style-type: none"> • GenAI exacerbates existing technology access inequities, including the availability of training, rising costs of tools, and availability of training and support for effective GenAI use (Fang & Zhou, 2025; Lake, 2024; Mimoudi, 2025; Suárez & García-Mariñoso, 2025; Liu & Wang, 2024). • GenAI development carries heavy, and unevenly distributed, costs, including environmental impacts caused by data centers and model training, the labor conditions of data workers, and disruptions to the job market (AWU-CWA & TechEquity, 2025; Furze, n.d.; Taft, 2025; Witt, 2025). • GenAI use raises unresolved legal and economic questions, including intellectual property and copyright disputes, the effects of automation on the workforce, and the impact of GenAI use on workload (Bondari, 2025; Brynolfsson et al., 2025; Cooper et al., 2025; Kinder et al., 2024; Lohwater, 2025; Ranganathan & Ye, 2026).

Appendix B: Building AI Literacy in Young Learners Activity Bank

Below are activities geared toward building AI literacy with young learners. Some are tech-free to honor early learning environments that should have minimal to no screentime. Activities that include using GenAI tools are focused on guided observation and teacher-facilitated demonstrations.

BUILDING KNOWLEDGE AND MINDSETS

AI or Not? sorting game

- Present images or short scenarios (a smart speaker, a pet dog, a calculator) and ask students to sort them into “AI” or “Not AI.” Discuss why each one landed where it did.
- For older learners, extend the sort into “AI” vs. “GenAI.” Ask, “Which of the AI examples retrieves and analyzes existing information? Which also uses that information to create something new?”

AI in My Life poster (older learners)

- Students identify AI-powered tools they encounter at home or school and label what each tool collects, analyzes, predicts, recommends, or creates.
- Follow with a gallery walk and discussion: “Which of these did you choose to use? Which just showed up in your life?”

BUILDING SAFE, ETHICAL, AND EFFECTIVE PRACTICES

Private or Public? sorting activity

- Use images and phrases (home address, birthday, pet’s name, favorite color) and have students sort them into safe versus unsafe to share online or with a “smart tool.”
- Discuss: “Who can see what you share online? What about with a smart tool? Does the tool remember it?”

Ask an Adult role play

- Act out scenarios where students encounter something confusing or uncomfortable with a tech tool (“This robot said something weird!”).
- Practice when to pause and find an adult for support.
- Discuss: “How do you know when something feels off? What do you do next?”

Who’s Missing? picture sort

- Generate a set of images of different roles: teachers, doctors, scientists, taxi drivers, etc.
- Have students identify who is represented in the pictures and who is missing.
- Talk about representation and inclusion.
- Discuss with students why the tool might show some people more than others.

What’s Honest? scenario sort

- Give students a variety of GenAI-use scenarios (using AI to check grammar, using it to write a poem, or creating a picture of someone).
- Have students sort them into “Honest,” “Not Honest,” or “It Depends,” and then explain their thinking.
- Discuss, “What makes something your work when you use GenAI to help create it?”

What’s a Useful Helper?

- Show photos or demos of technology tools students have encountered, such as a calculator, spell check, or a voice assistant.
- Discuss with students what each tool is able to do, when it is helpful, and when it might not be helpful to their learning.
- Then discuss when it isn’t appropriate to use the tool and to use their own brain, writing, or something else instead.

Appendix C:

Building AI Literacy in Secondary Learners

Activity Bank

Below are resources, activities and strategies for building GenAI literacy with early secondary learners. Note that this bank of activities and resources are suggestions for getting started. They can be done in concert with one another and applied to core academic standards and learning goals to more seamlessly integrate the development of GenAI literacy with mastery of academic and social-emotional knowledge and skills.

INTRODUCTION TO THE KNOWLEDGE, MINDSETS, AND PRACTICES

- [GenAI Literacy 101 for Students free course](#): our free, two-hour asynchronous course provides students 13 and older with a foundation of important knowledge, mindsets, and practices.

BUILDING KNOWLEDGE AND MINDSETS

Environmental impact discussion

- Use our [AI's Impact on the Environment classroom guide](#) to engage students in a conversation about the impact and what that means for them.
- This activity could be expanded into a more extensive research and debate assignment in alignment with persuasive writing and analysis standards.

Sycophancy test

- Have students make a clearly incorrect claim to a GenAI tool (e.g. “The moon is made of cheese” or “I’m going to stay up all night to study for an important exam”) and observe whether the tool pushes back or goes along with the claim.
- Discuss why the tool is designed to be agreeable, who made that choice, and why it might be a problem.

Same prompt, different tools

- In pairs or small groups, have students send the same prompt to three different tools (e.g. “Explain climate change to a seventh-grader”).
- Have students document the differences in content, tone, and accuracy.
- Discuss the differences and what they tell us about how the tools work, who made each tool, and how their choices impact the outputs.

BUILDING SAFE PRACTICES

Terms of Service speed read

- In pairs, give students 10 minutes to scan a GenAI tool’s terms of service and privacy policy with a simple checklist:
 - What data is collected?
 - Is it used for training?
 - Who can see it? Can you opt out?
- Then have students add their findings to posters around the room for a gallery walk.
- Discuss what was surprising, what they may have missed initially, and how the tool’s policy might impact how they would use it.

What Would You Type?

- Give students a list of information types such as relationship struggles, school assignment information, health concerns, opinions on current events, and family finances.
- Have them sort the information between “Would Share with GenAI Tool,” “Wouldn’t Share,” or “It Depends.” Discuss what influenced their decisions.
- Connect this activity to the “Terms of Service speed read” activity above if applicable.

Real Talk vs. Bot Talk

- Present a set of relatable personal dilemma scenarios.
- Have students work in trios and select one dilemma. One student discusses the dilemma and asks for support from a chatbot, while the other two students discuss it together with one offering support.
- Then have the trio compare the experiences.
- Discuss the difference in experiences among the class and the implications for when to turn to a human for support vs. a GenAI tool.

BUILDING ETHICAL PRACTICES

Understanding AI policies and guidelines

- Introduce district, school, or classroom guidelines for GenAI use.
- Have students write scenarios that would represent appropriate use and disclosure, as well as inappropriate use and disclosure.
- Have students share their scenarios with the class for discussion. For examples of inappropriate use, ask students, “What would need to be true for this scenario to be in line with our guidelines?”

BUILDING SAFE PRACTICES

Disclosure draft

- Give students realistic scenarios of how they might use GenAI to brainstorm ideas, get feedback, or revise written work.
- Have students write sample disclosure statements for the examples.
- Then have students share and reflect on what they would edit to accurately reflect their GenAI use.

Ethical prompt design challenge

- In teams, students create prompt guidelines for the ethical use of GenAI in different subject areas (e.g., journalism, art, science).
- Then have students present and defend their guidelines.

Uncovering Deepfakes

- Use our [Uncovering Deepfakes classroom guide](#) to engage students in understanding what deepfakes are and the impact they can have.
- Then discuss brief case studies where GenAI was used to harm, as in creating a deepfake image of a classmate, impersonating an adult in a message, or creating a false news report.
- Discuss the harm caused and what should be done to prevent these types of scenarios.

BUILDING EFFECTIVE PRACTICES

Bias audit

- Have students prompt a GenAI tool to develop images of people in different professions.
- Have them document patterns around gender, race, age, physical appearance, setting, etc.
- Then have students share their findings and discuss the implications of what GenAI tools consistently include vs. don't include.

The Five “S” prompting framework

- Use our [Five "S" prompting framework](#) to have students draft prompts related to a specific task for an upcoming project or assignment.
- Have students compare their prompts and the output with a partner and discuss the quality of their outputs, what they did differently, and how they can iterate on the prompt for a higher-quality product.

Student guide for AI Use flowchart and EVERY

- Introduce the [Student Guide for AI Use](#) and [EVERY](#) to students as ways to consistently apply the SEE Framework.
- Then have students evaluate a set of AI-use scenarios for schoolwork and personal use and discuss how they would apply each resource.

Student Self-Reflection checklist

- Use our [Student Self-Reflection checklist](#) after a project, unit, or grading period.
- Encourage students to discuss their results and next steps with a trusted adult or peer.

Appendix D: **Building AI Literacy in Adults Activity Bank**

Below are resources, activities, and strategies for building AI literacy with adult learners, including students over the age of 18, educators, administrators, caregivers, and community members. While the activities are designed with an education setting in mind, most can be adapted for a range of adult audiences. As with other sections of this framework, these suggestions work best when connected to participants' real contexts, roles, and goals.

INTRODUCTION TO THE KNOWLEDGE, MINDSETS, AND PRACTICES

- [An Essential Guide to AI for Educators](#): our free, two-hour asynchronous course provides educators with a foundation of important knowledge, mindsets, and practices.

BUILDING KNOWLEDGE AND MINDSETS

Myth vs. Fact

- Surface existing beliefs about GenAI from participants (or generate a list of common misconceptions).
- Have participants vote if the statement is a myth or fact. After each one, share what the evidence shows.
- Discuss the impact of understanding misconceptions vs. the truth of GenAI, its applications, and its impact in education.

Jagged Frontier mapping

- Introduce the concept that GenAI has a “jagged frontier” (Dell’Aqua et al., 2024).
- Provide a list of professional tasks (drafting an email, summarizing research, generating an image or map of a real place, solving a word puzzle, etc.) and have them predict what GenAI can handle well vs. handle poorly.
- Then test out some of the tasks live and reflect on how the understanding of the jagged frontier influences where they will use GenAI in work and life vs. where they won’t.

Community dialogue circle

- Facilitate a structured conversation in which participants from across the community surface their values, hopes, and concerns about GenAI.
- Then cluster the reflections and connect them to foundational core ideas (Table A1) such as capabilities and limitations, bias, training data, risk of overreliance, etc.
- Use paired discussion, whole group sharing, and a closing reflection to create a shared understanding of the community’s perspectives on GenAI and next steps for collaboration.

BUILDING SAFE PRACTICES

Terms of Service scavenger hunt

- Give pairs of participants a list of things to find in the terms of service and privacy policy of one of their preferred GenAI tools.
- Set a timer for 10 minutes and see how many items each pair can locate.
- Have pairs share what they found and discuss what surprised them, what didn’t surprise them, and how their understanding of the terms will impact their use of the tool.

What Should You Share?

- Have a list of information types relevant to the context of the audience (family information, student names, health concerns, staff performance data, business financial data, etc.).
- Have them sort the items into “Would Share,” “Wouldn’t Share,” or “It Depends.”
- Have participants compare their sorts and discuss what influenced their decisions.
- Connect the conversation to applicable data privacy policies

The AI Effect card game (from the Rithm Project)

- Using a set of [AI Effect cards](#), have participants sort the scenarios and discuss if GenAI is supporting or hindering human connection and why.
- As a whole group, discuss what scenarios surfaced tensions and differing perspectives and what that may mean for navigating the role GenAI may play in healthy, social, emotional development and human connection.

BUILDING ETHICAL PRACTICES

Understanding your organization's GenAI guidelines

- Introduce your organization's current guidelines for GenAI use (or review examples from peer organizations if your local guidance doesn't exist yet).
- Have participants capture questions of clarification.
- Then have them write two scenarios: one that is within the guidelines and one representing a gray area.
- Have them share in small groups and discuss what would need to be true for the ones in the gray area to be clearly within the guidelines.

Disclosure draft

- Give participants 2-3 scenarios in which they used GenAI to assist with a professional task, such as communication with families, developing discussion questions, or a policy memo.
- Have participants write sample disclosure statements that reflect what they contributed themselves vs. what the tool contributed.
- Compare statements, discuss differences, and align on what transparency should look like in their context.

When Is It OK?

- Present a range of scenarios for GenAI use relevant to the context and roles of the participants. Examples might include drafting feedback to students, generating a lesson plan on a current event, drafting accommodations in an IEP, screening job applications, or developing a meeting agenda.
- Have participants rate each scenario on a personal "ethical comfort scale."
- Then have participants share and compare in small groups discussing differences and what it reveals about choices for ethical practice.

Equity lens

- Provide a set of basic prompts (e.g., "Generate a picture of a kindergarten classroom in the spring," "write a story about a fourth-grade student struggling with a math lesson," or "draft a letter of recommendation for a high school senior that highlights [insert qualities]").
- Have participants generate outputs, then ask the tool itself to identify assumptions and biases in the outputs.
- Using what they find, have participants revise the prompts to invite more diverse perspectives and representation and compare the new outputs to the originals.
- Discuss the differences and what that shows about whose perspectives and knowledge are centered, both in the way we use the tools as well as in how the tools are designed.

BUILDING EFFECTIVE PRACTICES

Exploring effective prompting

- Share the [AI for Education Prompt Library](#) and [Five “S” Prompting Framework](#).
- Have participants explore a few sample prompts in the prompt library and review the prompting framework.
- Discuss what they notice about the structure of the prompts.
- Then have participants customize a prompt for the prompt library or write an original prompt using the Five “S” framework related to a task for their role.
- Ask them to iterate on the first output through evaluation and refinement.
- Discuss the key ideas they discovered about effective prompting and iteration when using GenAI tools.

Same task, different tools

- In small groups, have participants complete the same professional task with different GenAI tools (e.g. drafting a family newsletter, generating a discussion activity for a staff meeting, creating images for an upcoming project, vibe-coding an interactive for students).
- Have the groups compare the outputs from the various tools and make notes of differences in the outputs.
- Then have students capture what they observed for each tool in shared posters and do a gallery walk of findings.
- Discuss which tools were better at what and where different tools struggled. Have participants share the implications for their tool use moving forward.

Human vs. GenAI drafting

- Have participants choose a real task from their current work and spend five minutes capturing their own ideas for the task.
- Then have them use GenAI to develop the idea further using either text chat or voice.
- Have them discuss their outputs with a partner and share what components are reflective of their original thinking, creativity, and voice and what parts reflect the AI’s contribution.
- Discuss the impact on the quality of the output from starting with their own thinking first vs. having the AI generate the initial ideas.

References

Ahmmad, M., Shahzad, K., Iqbal, A., & Latif, M. (2025). Trap of Social Media Algorithms: A Systematic Review of Research on Filter Bubbles, Echo Chambers, and Their Impact on Youth. *Societies*, 15(11), 301. <https://doi.org/10.3390/soc15110301>

AI and the Global Economy: Unlocking Growth and Reshaping Work. (2025). LinkedIn. <https://economicgraph.linkedin.com/content/dam/me/economicgraph/en-us/PDF/ai-and-the-global-economy.pdf>

AI Literacy Lessons for Grades 6–12. (n.d.). Common Sense Media. Retrieved March 18, 2026, from <https://www.commonsense.org/education/collections/ai-literacy-lessons-for-grades-6-12>

aiEDU. (2025). AI readiness framework: What students, educators, and district leaders need to know. <https://www.aiedu.org/ai-readiness-framework>

Anderl, C., Klein, S. H., Sarigül, B., Schneider, F. M., Han, J., Fiedler, P. L., & Utz, S. (2024). Conversational presentation mode increases credibility judgements during information search with ChatGPT. *Scientific Reports*, 14(1), 17127. <https://doi.org/10.1038/s41598-024-67829-6>

Ash, A. M. (2025, June 25). Three in 10 Teachers Use AI Weekly, Saving Six Weeks a Year. *Gallup*. <https://news.gallup.com/poll/691967/three-teachers-weekly-saving-six-weeks-year.aspx>

AWU-CWA, & TechEquity. (2025). *Ghost Workers in the AI Machine: U.S. Data Workers Speak Out About Big Tech's Exploitation*. Communications Workers of America. <https://cwa-union.org/ghost-workers-ai-machine>

Bailey, H. (2025, November 14). AI “Trustwashing” Changes How Consumers Judge Credibility. *Tech Policy Press*. <https://techpolicy.press/ai-trustwashing-changes-how-consumers-judge-credibility>

Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, 33(1), 63. <https://doi.org/10.1007/s12525-023-00680-1>

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>

Belanger, A. (2025, November 7). *Oddest ChatGPT leaks yet: Cringey chat logs found in Google analytics tool*. *Ars Technica*. <https://arstechnica.com/tech-policy/2025/11/oddest-chatgpt-leaks-yet-cringey-chat-logs-found-in-google-analytics-tool/>

Berkowitz, J. S., Patock, J. R., Nawaz, A., Gonzalez-Hernandez, G., & Tatonetti, N. P. (2026). A crisis of overconfidence: Why confidence, not accuracy, is the real risk in clinical AI. *BioData Mining*, 19, 10. <https://doi.org/10.1186/s13040-026-00518-4>

Berson, I. R., Berson, M. J., & Luo, W. (2025). Innovating responsibly: Ethical considerations for AI in early childhood education. *AI, Brain and Child*, 1(1), 2. <https://doi.org/10.1007/s44436-025-00003-5>

Bickerstaff, A., Crouch, C. L., Factor, S., Fournier, J., Higgin, T., & Kosta, D. (2026). *Beyond the AI Inflection Point: Shaping the Future of Education*. AI for Education & Imagine Learning. <https://www.beyondtheaiinflectionpoint.com/>

Binder, M. (2026, January 21). *Researchers got Gemini AI to leak Google Calendar data, they claim*. Mashable. <https://mashable.com/article/google-gemini-ai-tricked-into-leaking-google-calendar-data>

Bondari, N. (2025, February 4). *AI, Copyright, and the Law: The Ongoing Battle Over Intellectual Property Rights*. USC. USC: Intellectual Property & Technology Law Society. <https://sites.usc.edu/iptls/2025/02/04/ai-copyright-and-the-law-the-ongoing-battle-over-intellectual-property-rights/>

Bradley, T. (n.d.). *One Prompt Can Bypass Every Major LLM's Safeguards*. Forbes. Retrieved March 11, 2026, from <https://www.forbes.com/sites/tonybradley/2025/04/24/one-prompt-can-bypass-every-major-llms-safeguards/>

Brynnfsson, E., Chandar, B., & Chen, R. (2025, November 13). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Stanford Digital Economy Lab. <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>

Burns, M., & Winthrop, R. (n.d.). *AI's future for students is in our hands*. Brookings. Retrieved March 12, 2026, from <https://www.brookings.edu/articles/ais-future-for-students-is-in-our-hands/>

Burns, M., Winthrop, R., Luther, N., Venetis, E., & Karim, R. (2026). *A new direction for students in an AI world: Prosper, prepare, protect*. Brookings. <https://www.brookings.edu/articles/a-new-direction-for-students-in-an-ai-world-prosper-prepare-protect/>

Capraro, V., Lentsch, A., Acemoglu, D., Akgun, S., Akhmedova, A., Bilancini, E., Bonnefon, J.-F., Brañas-Garza, P., Butera, L., Douglas, K. M., Everett, J. A. C., Gigerenzer, G., Greenhow, C., Hashimoto, D. A., Holt-Lunstad, J., Jetten, J., Johnson, S., Kunz, W. H., Longoni, C., ... Viale, R. (2024). The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS Nexus*, 3(6), pgae191. <https://doi.org/10.1093/pnasnexus/pgae191>

Chandra, K., Kleiman-Weiner, M., Ragan-Kelley, J., & Tenenbaum, J. B. (2026). *Sycophantic Chatbots Cause Delusional Spiraling, Even in Ideal Bayesians* (arXiv:2602.19141). arXiv. <https://doi.org/10.48550/arXiv.2602.19141>

Chang, D. H., Lin, M. P.-C., Hajian, S., & Wang, Q. Q. (2023). Educational Design Principles of Using AI Chatbot That Supports Self-Regulated Learning in Education: Goal Setting, Feedback, and Personalization. *Sustainability*, 15(17), 12921. <https://doi.org/10.3390/su151712921>

Chang, G., & Grant, H. (2026, January 23). When AI Amplifies the Biases of Its Users. *Harvard Business Review*. <https://hbr.org/2026/01/when-ai-amplifies-the-biases-of-its-users>

Chen Kulesa, A., Mission, M., Croft, M., & Wells, M. K. (2025). *Productive Struggle: How Artificial Intelligence Is Changing Learning, Effort, and Youth Development in Education*. Bellwether. <https://bellwether.org/publications/productive-struggle/>

Chen, L., Zhu, C., Soselia, D., Chen, J., Zhou, T., Goldstein, T., Huang, H., Shoeybi, M., & Catanzaro, B. (2024). *ODIN: Disentangled Reward Mitigates Hacking in RLHF* (arXiv:2402.07319). arXiv. <https://doi.org/10.48550/arXiv.2402.07319>

Chen, Z., Li, H., Liu, Z., Qu, H., & Xie, X. (2026). *From Passive Consumption to Active Interaction: Exploring Interactive LLM Scaffolding to Support Learning Engagement*. <https://doi.org/10.1145/3772363.3798931>

Cheng, J., Marone, M., Weller, O., Lawrie, D., Khashabi, D., & Durme, B. V. (2024). *Dated Data: Tracing Knowledge Cutoffs in Large Language Models* (arXiv:2403.12958). arXiv. <https://doi.org/10.48550/arXiv.2403.12958>.

Clegg, K.-A. (2025). Shoggoths, Sycophancy, Psychosis, Oh My: Rethinking Large Language Model Use and Safety. *Journal of Medical Internet Research*, 27, e87367. <https://doi.org/10.2196/87367>

Cooper, Z., Martens, B., Peukert, C., & Stocker, V. (n.d.). The Law and Economics of Generative AI and Copyright: A Primer to Core Challenges for Our Digital Future. *The Review of Network Economics*, 24(3), 161–176. <https://doi.org/10.1515/rne-2025-0039>

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., Besiroglu, T., & Owen, D. (2025). *The rising costs of training frontier AI models* (arXiv:2405.21015). arXiv. <https://doi.org/10.48550/arXiv.2405.21015>

Cuadra, A., Wang, M., Stein, L. A., Jung, M. F., Dell, N., Estrin, D., & Landay, J. A. (2024). The Illusion of Empathy? Notes on Displays of Emotion in Human-Computer Interaction. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, CHI '24*, 1–18. <https://doi.org/10.1145/3613904.3642336>

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality* (SSRN Scholarly Paper No. 4573321). Social Science Research Network. <https://doi.org/10.2139/ssrn.4573321>

Diliberti, M. K., Lake, R. J., & Weiner, S. R. (2025). *More Districts Are Training Teachers on Artificial Intelligence: Findings from the American School District Panel*. RAND. https://www.rand.org/pubs/research_reports/RRA956-31.html

Doctorow, C. (2024, October 30). *Pluralistic: AI's "human in the loop" isn't.*
<https://pluralistic.net/2024/10/30/a-neck-in-a-noose/#is-also-a-human-in-the-loop>

Doss, C. J., Bozick, R., Schwartz, H. L., Chu, L., Rainey, L. R., Woo, A., Reich, J., & Dukes, J. (2025). *AI Use in Schools Is Quickly Increasing but Guidance Lags Behind: Findings from the RAND Survey Panels.* RAND. https://www.rand.org/pubs/research_reports/RRA4180-1.html

Effective context engineering for AI agents. (2025, September 29). Anthropic.
<https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents>

Explaining Length Bias in LLM-Based Preference Evaluations. (n.d.). Retrieved March 11, 2026, from <https://arxiv.org/html/2407.01085v3>

Eynon, R. (2025, November 17). Expert Comment: Digital literacy on the curriculum is an opportunity for a more proactive response to AI in schools. *University of Oxford.*
<https://www.ox.ac.uk/news/2025-11-17-expert-comment-digital-literacy-curriculum-opportunity-more-proactive-response-ai>

Fabris, M. A., Roorda, D., & Longobardi, C. (2022). Editorial: Student-teacher relationship quality research: Past, present and future. *Frontiers in Education, 7.*
<https://doi.org/10.3389/feduc.2022.1049115>

Fang, L., & Zhou, X. (2025, April 2). *Will GenAI narrow or widen the digital divide in higher education?* SRHE Blog. <https://srheblog.com/2025/04/02/will-genai-narrow-or-widen-the-digital-divide-in-higher-education/>

Fanous, A., Goldberg, J., Agarwal, A., Lin, J., Zhou, A., Xu, S., Bikia, V., Daneshjou, R., & Koyejo, S. (2025). SycEval: Evaluating LLM Sycophancy. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 8*(1), 893–900. <https://doi.org/10.1609/aies.v8i1.36598>

Fesler, L., Martinez Claeys, J., Agnew, C., & Loeb, S. (2026). *The Evidence Base on AI in K-12: A 2026 Review.* AI Hub for Education of the SCALE Initiative.
<https://scale.stanford.edu/sites/default/files/The%20Evidence%20Base%20on%20AI%20in%20K-12%20Report.pdf>

Festerling, J., Siraj, I., & Malmberg, L.-E. (2022). Exploring children's exposure to voice assistants and their ontological conceptualizations of life and technology. *AI & Society, 1–28.*
<https://doi.org/10.1007/s00146-022-01555-3>

Freeman, J. (2025). *Student Generative AI Survey 2025* (Policy Note No. 61). Higher Education Policy Institute. <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/>

Friction by Design: A Framework for Centering Learning in the Age of AI. (2025). WestEd.
<https://www.wested.org/resource/friction-by-design/>

Furze, L. (n.d.). *Teaching AI Ethics: The Series*. Leon Furze. <https://leonfurze.com/ai-ethics/>

Gallegos, I. O., Rossi, R. A., Barrow, J., Tanjim, M. M., Kim, S., Deroncourt, F., Yu, T., Zhang, R., & Ahmed, N. K. (2024). Bias and Fairness in Large Language Models: A Survey. *Computational Linguistics*, 50(3), 1097–1179. https://doi.org/10.1162/coli_a_00524

Gerlich, M. (2025). AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies*, 15(1), 6. <https://doi.org/10.3390/soc15010006>

Green, B. (2022). *The Flaws of Policies Requiring Human Oversight of Government Algorithms* (SSRN Scholarly Paper No. 3921216). Social Science Research Network. <https://doi.org/10.2139/ssrn.3921216>

Grundmeier, R. W., Fiks, A. G., Jenssen, B. P., Proctor, S. N., Ferro, D. F., & Johnson, K. B. (2026). Generative Artificial Intelligence: Implications for Families and Pediatricians. *Pediatrics*, 157(4), e2025074912. <https://doi.org/10.1542/peds.2025-074912>

Guo, Y., Guo, M., Su, J., Yang, Z., Zhu, M., Li, H., Qiu, M., & Liu, S. S. (2024). *Bias in Large Language Models: Origin, Evaluation, and Mitigation* (arXiv:2411.10915). arXiv. <https://doi.org/10.48550/arXiv.2411.10915>

Health advisory: Artificial intelligence and adolescent well-being. (2025, June). American Psychological Association. <https://www.apa.org/topics/artificial-intelligence-machine-learning/health-advisory-ai-adolescent-well-being>

Heaven, W. D. (2026, January 7). LLMs contain a LOT of parameters. But what's a parameter? *MIT Technology Review*. <https://www.technologyreview.com/2026/01/07/1130795/what-even-is-a-parameter/>

Hoffmann, M. (2024, August 1). The Finalized EU Artificial Intelligence Act: Implications and Insights. *Center for Security and Emerging Technology*. <https://cset.georgetown.edu/article/the-finalized-eu-artificial-intelligence-act-implications-and-insights/>

How your data is used to improve model performance. (n.d.). OpenAI Help Center. Retrieved March 12, 2026, from <https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance>

Jacob, C., Kerrigan, P., & Bastos, M. (2025). The chat-chamber effect: Trusting the AI hallucination. *Big Data & Society*, 12(1), 20539517241306345. <https://doi.org/10.1177/20539517241306345>

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>

Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15(1), 17458. <https://doi.org/10.1038/s41598-025-97652-6>

Kim, H., Oh, S., & Lee, S. (2025a). *Mitigating Length Bias in RLHF through a Causal Lens*. <https://arxiv.org/html/2511.12573>

Kim, P., Xie, Y., & Yang, S. (2025b). *"I am here for you": How relational conversational AI appeals to adolescents, especially those who are socially and emotionally vulnerable* (arXiv:2512.15117). arXiv. <https://doi.org/10.48550/arXiv.2512.15117>

Kinder, M., de Souza Briggs, X., Muro, M., & Liu, S. (2024, October 10). Generative AI, the American worker, and the future of work. *Brookings*. <https://www.brookings.edu/articles/generative-ai-the-american-worker-and-the-future-of-work/>

King, J., Klyman, K., Capstick, E., Saade, T., & Hsieh, V. (2025). *User Privacy and Large Language Models: An Analysis of Frontier Developers' Privacy Policies* (arXiv:2509.05382). arXiv. <https://doi.org/10.48550/arXiv.2509.05382>

Lake, R. (2024, May 16). AI Is Coming to U.S. Classrooms, but Who Will Benefit? *Center on Reinventing Public Education*. <https://crpe.org/ai-is-coming-to-u-s-classrooms-but-who-will-benefit/>

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>

Lee, A., & Culver, M. (n.d.). *5 Principles for Prosocial AI: Guiding Young People Towards Human Connection*. The Rithm Project. Retrieved https://drive.google.com/file/d/1XRMG_ISSzLxR4v34O_6XrcEmH6GuEtib/view

Lee, H.-P. (Hank), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, CHI '25*, 1–22. <https://doi.org/10.1145/3706598.3713778>

Lehmann, M., Cornelius, P. B., & Sting, F. J. (2025). *AI Meets the Classroom: When Do Large Language Models Harm Learning?* (arXiv:2409.09047). arXiv. <https://doi.org/10.48550/arXiv.2409.09047>

Li, M., Chen, H., Wang, Y., Zhu, T., Zhang, W., Zhu, K., Wong, K.-F., & Wang, J. (2025). *Understanding and Mitigating the Bias Inheritance in LLM-based Data Augmentation on Downstream Tasks* (arXiv:2502.04419). arXiv. <https://doi.org/10.48550/arXiv.2502.04419>

Li, R.-N., Folk, D., Singh, A., Ungar, L., & Dunn, E. (2026). Is a random human peer better than a highly supportive chatbot in reducing loneliness over time? *Journal of Experimental Social Psychology*, 125, 104911. <https://doi.org/10.1016/j.jesp.2026.104911>

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns*, 4(7). <https://doi.org/10.1016/j.patter.2023.100779>

Lim, C. (2025). *DeBiasMe: De-biasing Human-AI Interactions with Metacognitive AIED (AI in Education) Interventions* (arXiv:2504.16770). arXiv. <https://doi.org/10.48550/arXiv.2504.16770>

Lira, B., Folk, D., Ungar, L., & Duckworth, A. L. (2026, January 28). How Gen Z Uses Gen AI—and Why It Worries Them. *Harvard Business Review*. <https://hbr.org/2026/01/how-gen-z-uses-gen-ai-and-why-it-worries-them>

Liu, S., Halder, K., Qi, Z., Xiao, W., Pappas, N., Htut, P. M., John, N. A., Benajiba, Y., & Roth, D. (2025). *Towards Long Context Hallucination Detection* (arXiv:2504.19457). arXiv. <https://doi.org/10.48550/arXiv.2504.19457>

Liu, Y., & Wang, H. (2024). *Who on Earth Is Using Generative AI?* [Policy Research Working Paper; 10870]. World Bank. <https://hdl.handle.net/10986/42071>

Lodge, J. M., & Loble, L. (2026). *Artificial intelligence, cognitive offloading and implications for education* [Report]. University of Technology Sydney. <https://doi.org/10.71741/4pyxmbnjq.31302475.v2>

Lohwater, T. (2025, November 18). *UC Berkeley experts discuss legal and economic questions about AI technologies*. UC Berkeley: College of Computing, Data Science, and Society. <https://cdss.berkeley.edu/news/uc-berkeley-experts-discuss-legal-and-economic-questions-about-ai-technologies>

Lopez-Lopez, E., Abels, C. M., Holford, D., Herzog, S. M., & Lewandowsky, S. (2025). Generative artificial intelligence–mediated confirmation bias in health information seeking. *Annals of the New York Academy of Sciences*, 1550(1), 23–36. <https://doi.org/10.1111/nyas.15413>

Loru, E., Nudo, J., Di Marco, N., Santirocchi, A., Atzeni, R., Cinelli, M., Cestari, V., Rossi-Arnaud, C., & Quattrociochi, W. (2025). The simulation of judgment in LLMs. *Proceedings of the National Academy of Sciences*, 122(42), e2518443122. <https://doi.org/10.1073/pnas.2518443122>

Machidon, O. M. (2025). Beyond Algorithethics: Addressing the Ethical and Anthropological Challenges of AI Recommender Systems. *Journal of Media Ethics*, 1–13. <https://doi.org/10.1080/23736992.2025.2584435>

Maheux, A. J., Akre-Bhide, S., Boeldt, D., Flannery, J. E., Richardson, Z., Burnell, K., Telzer, E. H., & Kollins, S. H. (2026). Generative Artificial Intelligence Applications Use Among US Youth. *JAMA Network Open*, 9(2), e2556631. <https://doi.org/10.1001/jamanetworkopen.2025.56631>

Malmqvist, L. (2025). Sycophancy in Large Language Models: Causes and Mitigations. In K. Arai (Ed.), *Intelligent Computing* (pp. 61–74). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-92611-2_5

Martineau, K. (2023, August 22). *What is retrieval-augmented generation (RAG)?* IBM Research. <https://research.ibm.com/blog/retrieval-augmented-generation-RAG>

Massenkoff, M., & McCrory, P. (2026). *Labor market impacts of AI: A new measure and early evidence*. Anthropic. <https://www.anthropic.com/research/labor-market-impacts>

McKendrick, J., & Thurai, A. (2022, September 15). AI Isn't Ready to Make Unsupervised Decisions. *Harvard Business Review*. <https://hbr.org/2022/09/ai-isnt-ready-to-make-unsupervised-decisions>

Meincke, L., Shapiro, D., Duckworth, A., Mollick, E. R., Mollick, L., & Cialdini, R. (2025). *Call Me A Jerk: Persuading AI to Comply with Objectionable Requests* (SSRN Scholarly Paper No. 5357179). Social Science Research Network. <https://doi.org/10.2139/ssrn.5357179>

Metz, C. (n.d.). Anthropic Agrees to Pay \$1.5 Billion to Settle Lawsuit With Book Authors. *The New York Times*. Retrieved March 11, 2026, from <https://www.nytimes.com/2025/09/05/technology/anthropic-settlement-copyright-ai.html>

Miao, F., & Cukurova, M. (2024). *AI competency framework for teachers*. UNESCO. <https://doi.org/10.54675/ZJTE2084>

Miao, F., Shiohira, K., & Lao, N. (2024). *AI competency framework for students*. UNSECO. <https://doi.org/10.54675/JKJB9835>

Mills, K., Ruiz, P., Lee, K., Coenraad, M., Fusco, J., Roschelle, J., & Weisgrau, J. (2024). *AI Literacy: A Framework to Understand, Evaluate, and Use Emerging Technology*. Digital Promise. <https://doi.org/10.51388/20.500.12265/218>

Mimoudi, A. (2025). Generative AI to bridge the educational divide: Personalized learning and challenges. *Social Sciences & Humanities Open*, 12, 102140. <https://doi.org/10.1016/j.ssaho.2025.102140>

Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus and Giroux.

Modi, A., Veerubhotla, A. S., Rysbek, A., Huber, A., Wiltshire, B., Veprek, B., Gillick, D., Kasenberg, D., Ahmed, D., Jurenka, I., Cohan, J., She, J., Wilkowski, J., Alarakyia, K., McKee, K. R., Wang, L., Kunesch, M., Schaekermann, M., Pîslar, M., ... Assael, Y. (2025). *LearnLM: Improving Gemini for Learning* (arXiv:2412.16429). arXiv. <https://doi.org/10.48550/arXiv.2412.16429>

Moore, J., Mehta, A., Agnew, W., Anthis, J. R., Louie, R., Mai, Y., Yin, P., Cheng, M., Paech, S. J., Klyman, K., Chancellor, S., Lin, E., Haber, N., & Ong, D. C. (2026, March 17). *Characterizing Delusional Spirals through Human-LLM Chat Logs*. arXiv.Org. <https://arxiv.org/abs/2603.16567v1>

Morrin, H., Nicholls, L., Levin, M., Yiend, J., Iyengar, U., DelGuidice, F., Bhattacharya, S., Tognin, S., MacCabe, J., Twumasi, R., Alderson-Day, B., & Pollak, T. A. (2026). Artificial intelligence-associated delusions and large language models: Risks, mechanisms of delusion co-creation, and safeguarding strategies. *The Lancet Psychiatry*. [https://doi.org/10.1016/S2215-0366\(25\)00396-7](https://doi.org/10.1016/S2215-0366(25)00396-7)

Munzer, T., Parga-Belinkie, J., Milkovich, L. M., Tomopoulos, S., Ajumobi, T., Cross, C., Gerwin, R., Madigan, S., Psych, R., & Council on Communications and Media. (2026). Digital Ecosystems, Children, and Adolescents: Policy Statement. *Pediatrics*, 157(2), e2025075320. <https://doi.org/10.1542/peds.2025-075320>

Myer, E., Nguyen, S., Edelson, L., & Mayer, J. (n.d.). *Tech Brief: AI Sycophancy & OpenAI*. *Georgetown Law*. Retrieved March 12, 2026, from <https://www.law.georgetown.edu/tech-institute/research-insights/insights/tech-brief-ai-sycophancy-openai-2/>

National Center for Education Statistics. (n.d.). *School Pulse Panel 2024-25*. Institute of Education Sciences. Retrieved March 18, 2026, from <https://nces.ed.gov/surveys/spp/default.asp>

National Scientific Council on the Developing Child. (2004). *Young Children Develop in an Environment of Relationships: Working Paper No. 1*. <https://developingchild.harvard.edu/resources/working-paper/wp1/>

Nikolopoulou, K. (2025). Child-centered integration of generative AI in early learning: Balancing promises and challenges. *AI, Brain and Child*, 1(1), 21. <https://doi.org/10.1007/s44436-025-00023-1>

Nolan, B. (2026, January 8). Google and Character.AI agree to settle lawsuits over teen suicides linked to AI chatbots. *Fortune*. <https://fortune.com/2026/01/08/google-character-ai-settle-lawsuits-teenage-child-suicides-chatbots/>

OECD. (2025). *Empowering learners for the age of AI: An AI literacy framework for primary and secondary education* (Review draft). OECD. <https://ailiteracyframework.org/>

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). *Training language models to follow instructions with human feedback* (arXiv:2203.02155). arXiv. <https://doi.org/10.48550/arXiv.2203.02155>

Peeperkorn, M., Kouwenhoven, T., Brown, D., & Jordanous, A. (2024). *Is Temperature the Creativity Parameter of Large Language Models?* (arXiv:2405.00492). arXiv. <https://doi.org/10.48550/arXiv.2405.00492>

Perkins, M., Roe, J., Vu, B. H., Postma, D., Hickerson, D., McGaughran, J., & Khuat, H. Q. (2024). Simple techniques to bypass GenAI text detectors: Implications for inclusive education. *International Journal of Educational Technology in Higher Education*, 21(1), 53. <https://doi.org/10.1186/s41239-024-00487-w>.

Perry, A. (2023). AI will never convey the essence of human empathy. *Nature Human Behaviour*, 7(11), 1808–1809. <https://doi.org/10.1038/s41562-023-01675-w>

Pew Research Center. (2025). *Teens, Social Media and AI Chatbots 2025*. <https://www.pewresearch.org/internet/2025/12/09/teens-social-media-and-ai-chatbots-2025/>

Pew Research Center. (2026). *How Teens Use and View AI*. <https://www.pewresearch.org/internet/2026/02/24/how-teens-use-and-view-ai/>

Powers-O'Brien, A., & Rochman, A. (2025). *Generative AI in Uncertain Times: How Teens are Navigating a New Digital Frontier*. Family Online Safety Institute. <https://fosi.org/research/generative-ai-in-uncertain-times-how-teens-are-navigating-a-new-digital-frontier/>

Prothero, A. (2024, February 19). What Is Age-Appropriate Use of AI? 4 Developmental Stages to Know About. *Education Week*. <https://www.edweek.org/technology/what-is-age-appropriate-use-of-ai-4-developmental-stages-to-know-about/2024/02>

Prothero, A., & Vilcarino, J. (2025, October 20). How School Districts Are Crafting AI Policy on the Fly. *Education Week*. <https://www.edweek.org/technology/how-school-districts-are-crafting-ai-policy-on-the-fly/2025/10>

Quattrociochi, W., Capraro, V., & Perc, M. (2025). *Epistemological Fault Lines Between Human and Artificial Intelligence* (arXiv:2512.19466). arXiv. <https://doi.org/10.48550/arXiv.2512.19466>

Ranganathan, A., & Ye, X. M. (2026, February 9). AI Doesn't Reduce Work—It Intensifies It. *Harvard Business Review*. <https://hbr.org/2026/02/ai-doesnt-reduce-work-it-intensifies-it>

Reich, J. (2025, October 3). What past education technology failures can teach us about the future of AI in schools. *The Conversation*. <https://doi.org/10.64628/AAI.vvdntn96m>

Riemer, K., & Peter, S. (2025, March 6). AI doesn't really 'learn' – and knowing why will help you use it more responsibly. *The Conversation*. <https://doi.org/10.64628/AA.we96s45cj>

Robb, M., & Mann, S. (2025). *Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions*. Common Sense Media. <https://www.common sense media.org/research/talk-trust-and-trade-offs-how-and-why-teens-use-ai-companions>

Robbins, S., & Blundell, I. (2025). Losing Our Voice? Generative AI and the Degradation of Human Expression. *Minds and Machines*, 36(1), 2. <https://doi.org/10.1007/s11023-025-09757-6>

Saito, K., Wachi, A., Wataoka, K., & Akimoto, Y. (2023). *Verbosity Bias in Preference Labeling by Large Language Models* (arXiv:2310.10076). arXiv. <https://doi.org/10.48550/arXiv.2310.10076>

Salvaggio, E. (2025, June 17). The Black Box Myth: What the Industry Pretends Not to Know About AI. *Tech Policy Press*. <https://techpolicy.press/the-black-box-myth-what-the-industry-pretends-not-to-know-about-ai>

Sanwoolu, O. D. (2025). Kantian Deontology for AI: Alignment Without Moral Agency. *AI and Ethics*, 5, 5425–5437. <https://doi.org/10.1007/s43681-025-00784-8>

Schechner, S., & Kessler, S. (2025, August 7). ‘I Feel Like I’m Going Crazy’: ChatGPT Fuels Delusional Spirals. *The Wall Street Journal*. <https://www.wsj.com/tech/ai/i-feel-like-im-going-crazy-chatgpt-fuels-delusional-spirals-ae5a51fc>

Sengar, S. S., Hasan, A. B., Kumar, S., & Carroll, F. (2025). Generative artificial intelligence: A systematic review and applications. *Multimedia Tools and Applications*, 84(21), 23661–23700. <https://doi.org/10.1007/s11042-024-20016-1>

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askill, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2025). *Towards Understanding Sycophancy in Language Models* (arXiv:2310.13548). arXiv. <https://doi.org/10.48550/arXiv.2310.13548>

Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, CHI '24*, 1–17. <https://doi.org/10.1145/3613904.3642459>

Shaw, S. D., & Nave, G. (2026). *Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender* (SSRN Scholarly Paper No. 6097646). Social Science Research Network. <https://doi.org/10.2139/ssrn.6097646>

Shen, J. H., & Tamkin, A. (2026). *How AI Impacts Skill Formation* (arXiv:2601.20245). arXiv. <https://doi.org/10.48550/arXiv.2601.20245>

Shin, D., & Shin, E. Y. (2026). Cascading falsehoods: Mapping the diffusion of misinformation in algorithmic environments. *AI & SOCIETY*, 41(2), 1435–1452. <https://doi.org/10.1007/s00146-025-02575-5>

- Singh, A., Taneja, K., Guan, Z., & Ghosh, A. (2025). Protecting Human Cognition in the Age of AI (arXiv:2502.12447). arXiv. <https://doi.org/10.48550/arXiv.2502.12447>
- Singhal, P., Goyal, T., Xu, J., & Durrett, G. (2024). *A Long Way to Go: Investigating Length Correlations in RLHF* (arXiv:2310.03716). arXiv. <https://doi.org/10.48550/arXiv.2310.03716>
- Sourati, Z., S. Ziabari, A., & Dehghani, M. (2026). The homogenizing effect of large language models on human expression and thought. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2026.01.003>
- Suárez, D., & García-Mariñoso, B. (2025). On the verge of a digital divide in the use of generative AI? *Telecommunications Policy*, 49(7), 102997. <https://doi.org/10.1016/j.telpol.2025.102997>
- Surgo Health, & The JED Foundation. (2026). *Youth Mental Health in the AI Era: How GenAI Enters Help-Seeking Pathways*. <https://www.trackyouthmentalhealth.com/reports/youth-mental-health-in-the-ai-era-how-genai-enters-help-seeking-pathways>
- Suzgun, M., Gur, T., Bianchi, F., Ho, D. E., Icard, T., Jurafsky, D., & Zou, J. (2025). Language models cannot reliably distinguish belief from knowledge and fact. *Nature Machine Intelligence*, 7(11), 1780–1790. <https://doi.org/10.1038/s42256-025-01113-8>
- Sycophancy in GPT-4o: What happened and what we're doing about it.* (2025, April 29). OpenAI. <https://openai.com/index/sycophancy-in-gpt-4o/>
- Tadimalla, S. Y., & Maher, M. L. (2025). AI literacy as a core component of AI education. *AI Magazine*, 46(2), e70007. <https://doi.org/10.1002/aaai.70007>
- Taft, M. (2025, December 16). How Much Water Do AI Data Centers Really Use? *Undark Magazine*. <https://undark.org/2025/12/16/ai-data-centers-water/>
- Tao, S., Lan, M., Wang, M., & Li, H. (2026). Potential risks of generative artificial intelligence integration into K-12 education: A scoping review. *Computers and Education: Artificial Intelligence*, 10, 100561. <https://doi.org/10.1016/j.caeai.2026.100561>
- The Children and Screens Guide for Early Child Development and Media Use: Infants and Children Ages 0-5.* (n.d.). Children and Screens. Retrieved March 12, 2026, from <https://www.childrenandscreens.org/learn-explore/research/the-children-and-screens-guide-for-early-child-development-and-media-use-infants-and-children-ages-0-5/>
- The Ongoing Fight for A.I. Protections Makes Waves on Capitol Hill and Beyond.* (2024, August 16). SAG-AFTRA. <https://www.sagaftra.org/ongoing-fight-ai-protections-makes-waves-capitol-hill-and-beyond>

Trujillo, M. (2024, July 31). Texas Wins \$1.4 Billion Biometric Settlement Against Meta. It Would Have Happened Sooner With Consumer Enforcement. *Electronic Frontier Foundation*. <https://www.eff.org/deeplinks/2024/07/texas-wins-14-billion-biometric-settlement-against-meta-it-would-have-happened>

Tully, S. M., Longoni, C., & Appel, G. (2025). Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity. *Journal of Marketing*, 89(5), 1–20. <https://doi.org/10.1177/00222429251314491>

Ulrich, C. M., Oh, O., Bin You, S., Topaz, M., Rahemi, Z., Stokes, L., Pruinelli, L., Demiris, G., & Flatley Brennan, P. (2026). What Does Moral Agency Mean for Nurses in the Era of Artificial Intelligence? *The Hastings Center Report*, 56(1), 18–23. <https://doi.org/10.1002/hast.70030>

Understanding the Impacts of Generative AI Use on Children: Recommendations. (2025). The Alan Turing Institute. https://www.turing.ac.uk/sites/default/files/2025-05/understanding_the_impacts_of_generative_ai_use_on_children_-_recommendations.pdf

Understanding tokens. (2025, May 29). Microsoft Learn. <https://learn.microsoft.com/en-us/dotnet/ai/conceptual/understanding-tokens>

Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293–2303. <https://doi.org/10.1038/s41562-024-02024-1>

Varughese, J. (2026, February 20). *What Is LLM Training?* IBM. <https://www.ibm.com/think/topics/llm-training>

Vassel, F.-M. (2024, October 3). *How Harmful Are AI's Biases on Diverse Student Populations?* (P. Kannan, Interviewer). Stanford University: Human-Centered Artificial Intelligence. <https://hai.stanford.edu/news/how-harmful-are-ais-biases-on-diverse-student-populations>

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). *Attention Is All You Need* (arXiv:1706.03762). arXiv. <https://doi.org/10.48550/arXiv.1706.03762>

Weizenbaum, J. (1966). ELIZA—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45. <https://doi.org/10.1145/365153.365168>

Wiggers, K. (2024, May 4). Why RAG won't solve generative AI's hallucination problem. *TechCrunch*. <https://techcrunch.com/2024/05/04/why-rag-wont-solve-generative-ais-hallucination-problem/>

Willison, S. (2024, May 29). Training is not the same as chatting: ChatGPT and other LLMs don't remember everything you say. *Simon Willison's Weblog*.
<https://simonwillison.net/2024/May/29/training-not-chatting/>

Witt, S. (n.d.). Inside the Data Centers That Train A.I. and Drain the Electrical Grid. *The New Yorker*. Retrieved March 12, 2026, from <https://www.newyorker.com/magazine/2025/11/03/inside-the-data-centers-that-train-ai-and-drain-the-electrical-grid>

Xu, Y. (Guest Expert), Perry, K. (Host). (2025, April). *AI's Impact on Children's Social and Cognitive Development* (No. 013) [Broadcast]. Children and Screens.
<https://www.childrenandscreens.org/learn-explore/research/ais-impact-on-childrens-social-and-cognitive-development-ying-xu-phd/>

Zeff, M. (2025, June 2). How AI chatbots keep you chatting. *TechCrunch*.
<https://techcrunch.com/2025/06/02/how-ai-chatbots-keep-you-chatting/>

Zhao, C. (2025, October 28). AI hallucinates because it's trained to fake answers it doesn't know. *Science*. <https://www.science.org/content/article/ai-hallucinates-because-it-s-trained-fake-answers-it-doesn-t-know>

Generative AI Disclosure Statement

This document was developed with support from Claude to assist in content feedback and refinement. All AI-generated contributions were reviewed, edited, and adapted to align with the document's goals and intended tone. Final decisions about structure, messaging, and content were made by human authors to align with pedagogical intent and support responsible AI use across diverse learning contexts and audiences.

For inquiries, permissions, or collaboration requests, contact: info@aiforeducation.io
Website: <https://www.aiforeducation.io/>

AI for Education

aiforeducation.io

AI for Education | SEE FRAMEWORK 2026