AI Bytes

Making Sense of Artificial Intelligence in Get SET (Skills, Education and Training)

A Contact North | Contact Nord and Literacy Link South Central publication

e-Channel
Apprentissage en ligne

Literacy Link
South Central

Get SET
Skills, Education and Training

Welcome to *AI Bytes*, your curated guide to making sense of AI in Get SET (Skills, Education and Training).

These days, a learner can type one sentence and voila! … an instant paragraph appears. It makes this question feel louder than ever: **"If AI can do this instantly, why bother learning?"**

It's a fair question. But remember Joey Tribbiani from Friends — The One with the Thesaurus (2:53)?

When he used a thesaurus to "sound smart," he turned "*They are warm nice people with big hearts*" into "*They are humid pre-possessing homo sapiens with full-sized aortic pumps*."

AI can be a bit like Joey using a thesaurus to write a recommendation letter: polished on the surface but not always making sense underneath. In the episode, Joey swaps every word for a fancier one and ends up with a sentence that sounds impressive but collapses under basic scrutiny.

AI can sound the same: fluent, confident and occasionally way off base unless a human checks the meaning.

Learning still matters because memorizing is only the starting point; the real work is meaning, judgment and voice. The human ability to slow down with AI and keep your own voice in the mix is critical.

Pour yourself a coffee and join us. In this issue:

- What you need to know about large language models (LLMs)
- AI can generate text but cannot replace learning
- From theory to practice: Experimenting with AI intentionally

Learning matters because the real work begins where AI ends — judgment, depth, meaning.

Ontario
This Employment Ontario program is funded by the Government of Ontario

Meet the *AI Bytes* team

Carolina Cohoon is an EdTech Consultant at Literacy Link South Central. Her professional background encompasses education and rehabilitation, with a passion for inclusion and accessibility. Carolina is dedicated to designing learning experiences that celebrate and embrace diversity. Her interest in AI is fueled by her enthusiasm for innovation, knowledge sharing, enhancing accessibility, and improving the learning experience through personalized learning adaptations that AI can offer within the framework of Universal Design for Learning (UDL). Carolina is certified in ChatGPT through the Blockchain Council and has completed Ivey School of Business training in Accelerating Leadership through AI, demonstrating her dedication to advancing inclusive, technology-enhanced education that drives meaningful transformation in learning.

Jeremy Marks works for Literacy Link South Central as a project manager and Ed. Tech researcher. He also teaches Communications in the ACE program at Fanshawe College, London. Jeremy is a certified teacher/trainer of adults and has taught learners in public and secondary schools, colleges and universities, in Canada and the U.S., since 2002. His fascination with AI comes from his longstanding passion for educational theory, cognitive philosophy, and yes, science fiction. In his spare time, Jeremy writes books (poetry & fiction) and is a published author.

* This bulletin is edited by Contact North | Contact Nord. Generative AI tools supported aspects of ideation, formatting and image creation in this work. Original ideas, translation, knowledge, research and connections were developed by the authors. This disclosure reflects our commitment to transparency, intellectual integrity and responsible use of emerging technologies.

What you need to know about Large Language Models (LLM)

A custom definition for Get SET:

Large language models, or LLMs, are pattern-based systems trained to generate the most likely next token in a sequence. Much of the fluency, apparent reasoning and confident-sounding output we see comes from that training process.

LLMs are one type of generative AI, but they are the type learners are most likely to use for brainstorming, outlining, explaining and summarizing. Generative AI is the broader family, including text, image, audio, video, code and multimodal systems.

Inspired by Andreas Stöffelbauer's How Large Language Models Work: From Zero to ChatGPT*, Data Science at Microsoft (Medium, 2023), with additional framing for AI literacy in adult education. Portions of the explanation of GenAI capabilities are informed by the University of Calgary's AI Literacy guide, specifically the section What GenAI Can Do (brainstorming, outlining, explaining, supporting writing and synthesizing information).*

Mainstream definition of LLM

An LLM is a large neural network with billions of parameters. Its core task is deceptively simple: predict the next token in a sequence. [Tokens are not always whole words]; they may be full words, parts of words or punctuation.

AI cannot fully substitute for the meaning-making that comes from lived experience, dialogue and reflection. It learns patterns at enormous scale, and what feels like understanding emerges from learned statistical patterns and internal representations, not from lived experience or grounded comprehension.

How LLMs learn

Pre training: The model processes massive amounts of text and repeatedly predicts the next token. These prediction cycles form the statistical patterns it uses for grammar, structure and common language behaviours.

Instruction fine tuning: Curated examples show the model how to follow instructions. This shifts it from simply continuing text to producing responses that match patterns such as summaries, explanations and steps.

Reinforcement Learning from Human Feedback (RLHF): Human raters compare multiple model outputs and rank them. Those rankings are used to train a reward model that approximates human preferences. The language model is then optimized, often through reinforcement learning, to increase the likelihood of responses that score higher on that reward model, shifting its behaviour toward patterns humans judged as more helpful, safe and appropriate.

What is temperature?

Temperature is one of the simplest and most powerful settings that shape how AI models respond. It determines whether an AI sounds factual and precise or creative and expressive.

Think of it as a slider that balances predictability against surprise:

- Low temperatures lead to increased accuracy and stability. This is exactly what you want for coding, math, or summarizing data.
- Medium temperatures deliver a natural balance and fluency. This is the natural back-and-forth you get from everyday chatbots.
- High temperatures spark creativity and risk-taking. This setting is best for brainstorming or storytelling.

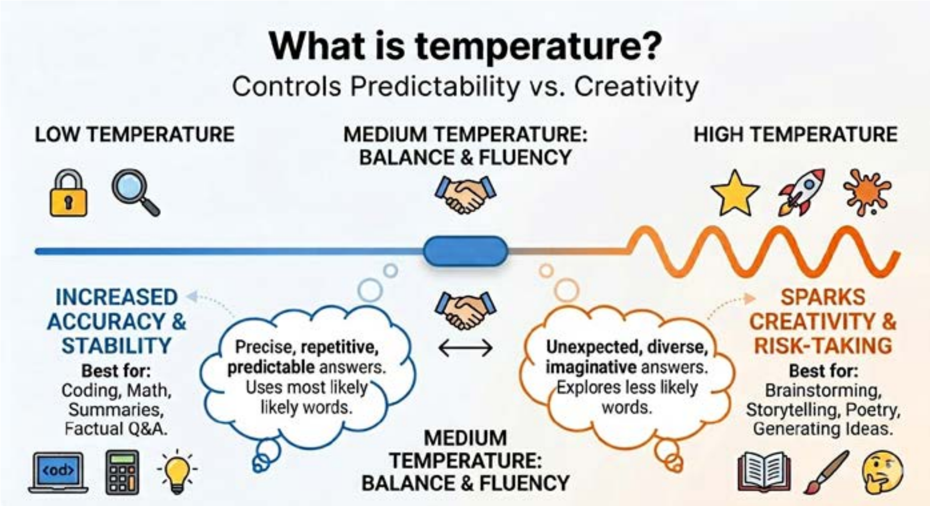

What's the recommended temperature?

The short answer is: it depends on your goals. As the team at [Watercrawl] highlights, there's no one-size-fits-all setting. The right choice depends on your task, whether that's generating code, writing summaries, or crafting creative stories. Experimenting is the best way to find the right balance between consistency and creativity. One important caveat, though: [a low temperature ensures consistency, not accuracy]. If the model lacks the information, lowering the temperature simply guarantees the same error will be reproduced with every run.

Want to try it yourself? A sandbox like [Google AI Studio]() lets you move the slider in real time and see the difference firsthand. Explore exactly how shifting the temperature impacts the tone, logic, and creativity of your chats!

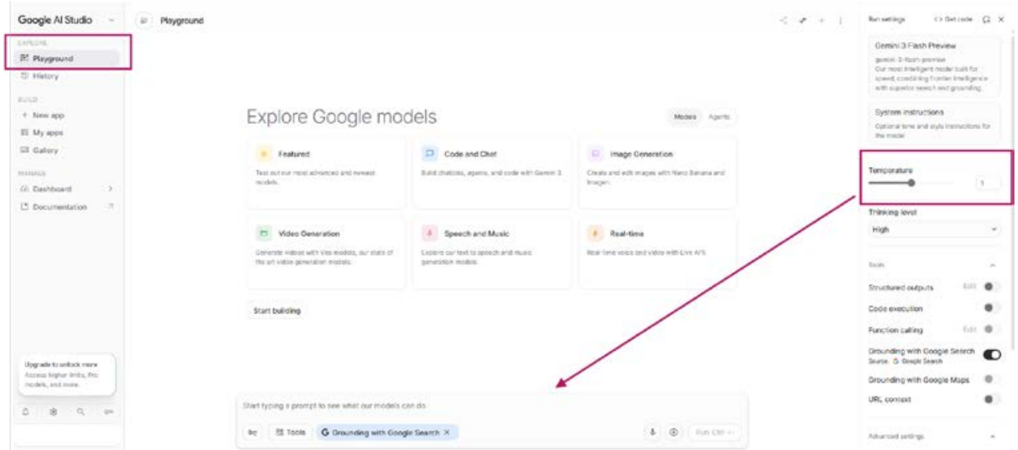

LLMs seem so capable but they have limitations

LLMs can generate fluent text, summarize material, draft ideas, explain concepts and respond in ways that sound polished and confident. That is part of why they can feel so convincing, even when the content still needs verification.

LLMs, through large scale pattern training, can:

Capability	Limitations
Produces paragraphs	No guarantee of proper structure or correction information.
Summarizes documents	Needs verification.
Translates languages	Quality isn't equal. English dominates most training sets, and many languages are still under-represented.
Generates and explains code	The quality depends on the patterns in its training data — and there's plenty of low-quality code online. It's great at formatting but the output still needs human review.
Formats text	LLMs are strong at making text look neat, but they don't maintain structure reliably — formatting can drift, collapse, or look polished while still being logically or technically wrong.
Answers many factual questions	Can be incomplete, wrong, or outdated.
Mimics different writing styles	Although good at the surface level, style can drift, nuance gets flattened, and the writing can feel off to a fluent reader.

A model may produce a response that looks insightful because it has learned how insight is typically expressed in its training data. The apparent skill comes from pattern learning at scale, not from human-style understanding.

The following graphic outlines some capabilities and [limitations when using LLMs]():

Additional capabilities and limitations can be found at: [AI Literacy - Artificial Intelligence - Library at University of Calgary]()

Why LLMs make mistakes

LLMs are statistical text generation systems. They are trained to continue a sequence of tokens based on patterns in their training data, not to reliably verify facts or track truth.

Hallucinations

Human writing often sounds confident, so LLMs learn that style. As a result, they may generate fluent sounding language even when the information is unreliable, sometimes leading to fabricated details, invented citations or distorted sources.

These errors, often called hallucinations, reflect the model's training objective and remain an active research area. Hallucinations are also more likely when the model lacks access to reliable external sources or when prompts require precise factual recall.

Bias

LLM outputs can also reproduce social biases that are present in the data they were trained on. For example, some models show stronger associations between certain roles or traits and particular genders, races or age groups, which can shape how people and communities are described. These biases are not always obvious, but they can influence how certain experiences, accents or life paths appear in the model's responses. These patterns emerge because training data reflects historical and social inequalities that the model can learn and reproduce unless actively mitigated.

Instruction tuning, which teaches LLMs to follow explicit instructions, can reduce biases but it is not yet clear whether they can be fully eliminated.

LLM limits

Even the most advanced models have important limitations on their own:

- They are not reliable for human-like understanding
- They are not dependable truth verifiers
- They do not reason with human accountability*
- They do not form beliefs, intentions or values
- They do not make ethical or moral judgments in the human sense
- They cannot guarantee that a citation exists or says what the model claims it says

Some AI systems combine generation with retrieval (e.g., searching documents or the web). These systems can reduce hallucinations by grounding responses in sources, but they still require verification because they may misinterpret or selectively use that information.

**They can simulate reasoning and solve many problems, but this process is not grounded in understanding or accountability in the human sense.*

Tokenization

One of the most effective ways to demystify LLMs is to understand tokens, the actual units an LLM reads and writes.

An LLM does not necessarily process whole words. It breaks text into tokens: chunks that might be a full word, a partial word, a punctuation mark or even a single character. The word "unhappiness", for example, might be split into three tokens: un · happi · ness.

Here's a breakdown:

- **"un"** is a common prefix that appears as its own token
- **"happi"** is a frequent stem (from "happy," "happiness," "happily")
- **"ness"** is a common suffix that is also its own token

Every time an LLM generates text, it weighs the probabilities of possible next tokens and then repeats the process again, one token at a time.

As the model repeats this step, token by token, the output emerges. The process is probabilistic rather than conceptual, shaped by learned patterns and context rather than lived understanding.

Understanding these concepts promotes AI literacy. When a response sounds authoritative, they can ask: Is this the most likely phrasing or the most accurate one? LLMs are optimized for likely text rather than truth, so they may answer the first more reliably than the second.

The big idea

AI can feel like an intimidating technical shift, but educators don't need to be tech experts to lead this conversation. We are all learning in real time.

To use AI effectively, learners need one core insight: AI can generate strong drafts, but it does not verify facts on its own. When we talk openly about where these systems succeed and where they fail, the technology becomes more manageable. That matters because relying on AI without oversight can limit a learner's options and growth in the workplace. Every AI response should be treated as a starting point rather than a final answer.

"Clearly these LLMs are proving to be very useful and show impressive knowledge and reasoning capabilities, and maybe even show some sparks of general intelligence. But whether or to what extent that resembles human intelligence is still to be determined" (Stöffelbauer, 2023, conclusion section, para. 2).

What this means in practice

Knowing that AI can't replace learning is only the starting point. The bigger question is how to make that limitation visible to learners. One way is to help them examine how AI works, test its outputs critically and notice where their own judgment still matters.

From there, AI can be introduced as a support for learning rather than a substitute for it. Learners can explore tokenization and prediction, analyze AI drafts for bias or missing context and document their collaboration with AI through a disclosure or collaborator statement. These practices make learning more visible and reinforce the role of human effort.

In this sense, AI literacy is not just about using tools well. It is about knowing when to question them, how to verify them and why the learner's own thinking remains central. Below are some tips to introduce LLM concepts:

01

Discuss tokenization and prediction

Before learners can use AI critically, they need to understand what it actually is. Start with tokenization: AI doesn't read sentences the way we do, it breaks text into fragments called tokens and predicts the most probable next one until a response is complete.

Activity 1: Use the Tokenizer Playground to show learners how different sentences get broken up. Type a few examples together and watch how the text splits into colored segments. Note that the colors themselves such as purple, yellow, red carry no meaning. They simply cycle through a palette to make it easy to see where one token ends and the next begins.

Activity 2: Type a prompt like "*The capital of France is*" and ask learners to predict what comes next. Then show how the model completes it. Explain it will say "Paris" not because it "knows" the answer, but because "Paris" is statistically the most likely continuation of that phrase. Follow that with a trickier prompt: ask the model to describe a made-up animal like an "*Octoviotopus*" and discuss the output This illustrates how the model generates plausible-sounding text even when there's nothing real to draw from.

02

Make AI outputs a subject of study

Work with learners to analyze AI-generated drafts. Ask them to identify where the output relies on bias, hallucinates details, misses local context or flattens nuance.

Learners who can articulate why a specific AI draft is weak are demonstrating the kind of disciplinary reasoning that generic AI does not reliably perform.

Activity 3: Have the machine draft a quick overview of native flowers in Southwestern Ontario. Once generated, invite your learners to audit the output: "Whose voices and perspectives are centered here, and whose are completely left out?" This reveals a vital truth about AI literacy: LLMs are never neutral or universal. They simply mirror the statistical averages of the dominant data and people that trained them.

03

Build a collaborator statement

Invite learners to review existing AI statements and disclosures, then co create a collaborator statement to accompany any work where AI was used. This statement should outline which tool was used, what prompts were given, which parts were kept or changed and what the learner contributed in terms of accuracy, voice, judgment and original thinking.

This practice helps make AI use more transparent and accountable, while helping learners articulate their own intellectual contribution. That articulation is itself a learning process, and a transferable workplace skill.

For adult learners and instructors, responsible AI use means being transparent about when and how AI was used in assignments, communications and professional work.

Activity 4: We created an Ethic Statement generator, follow this link to learn how to use it in your classroom: [Exploring Google Opal].

04

Build artifacts: Workflows, prompt guides and comparison tables

Shift the deliverable from product to process. Encourage learners to build and share prompt guides that document what worked and why, comparison tables showing AI output versus their revised version, or documented workflows that show AI as one step in a larger thinking process.

These artifacts make the process more visible, capturing the reasoning, judgment and revision. They are also more AI-resistant than essays on generic topics.

Activity 5 (Advanced and No-Code): Open the [Ethics Statement app] along with its backend editor. Ask learners to observe the visual workflow. This shows them exactly how natural language instructions generate the final statement. Alternatively, invite learners to pitch their own unique app ideas, and together, prompt Opal to build the app from scratch. It could be something as simple as "I want to build a single-page 'Budget Grocery Planner' app that tracks my grocery list alongside a running total cost, helping me plan and stay under a set budget limit of $120 per week".

05

Explain AI sycophancy

Work with learners to analyze how LLMs sometimes prioritize helpfulness over accuracy. Ask them to identify where the AI is merely echoing their own prompt back to them rather than providing a rigorous or objective critique.

Teaching tips:

- Have learners provide the AI with a strong, even incorrect, opinion. Ask them to find the exact sentence where the AI appears to stop challenging the user and starts agreeing too readily.
- Look for overly apologetic or validating language (e.g., "*You've raised an important point that is often overlooked*...") and discuss how this language masks a lack of substance.
- Analyze where the AI ignores local facts or specific cultural nuances because the user's prompt suggested a different narrative.

Building quality control with LLM

Since LLMs default to "completion" rather than "verification," learners should build their own quality control process. A few tips:

- Cross-reference answers. Copy a prompt into a second, different LLM. If one model says something different from another, that can be a sign that the information should be checked further.
- Do not read only the AI's response. Read other sources in separate browser tabs.
- Confirm any bold or surprising claim with at least three independent, authoritative sources such as government websites, academic journals or reputable news outlets.
- Use fail-safe prompts. Ask the model to say "I don't know" rather than guess if it is not certain.
- Define boundaries. For example: "Only use information from this year" or "Only use information from this organization."
- Ground the model in the text or links you want analyzed.

You might also invite learners to notice how the model's tone changes when they adjust temperature or ask for "more confident" or "more creative" answers. This helps them see that style and confidence are partly a design choice, not a measure of truth.

AI can generate text, but it cannot replace learning

Remember Cole's Notes?

If you attended school in Canada in the late 20th century, these yellow-and-black booklets may be familiar. Cole's Notes offered *Hamlet*, for example, in a much shorter form, but the notes did not replace the reader's own drafting, interpreting or writing. AI is different because it can generate the draft for you. Even if the output still requires editing, a tool that produces the draft can shift from a supplement to a shortcut around thinking, both being forms of cognitive offloading.

The following comic illustrates this shift:

A field experiment reported that unrestricted GPT-4 support improved short-term math practice performance, but when AI access was removed, those students later performed worse than students who had never used AI. A safeguarded version of the tutor improved practice performance while largely avoiding that learning loss.

The reported performance gains were substantial:

48% improvement
for students using the unrestricted GPT 4 tutor

127% improvement
for those using the safeguarded version

But the real insight came later. When AI access was removed, students who had used the unrestricted version performed worse than students who had never used AI at all. In other words, without guardrails, some learners may have leaned on GPT-4 during practice, and their independent performance later dropped (17%).

The safeguarded tutor avoided this collapse. Its design nudged students to stay in the cognitive process rather than outsource it. The authors' conclusion suggests: **AI can boost performance, but design choices determine whether it strengthens or weakens learning.**

John Nosta (*Psychology Today*, 2026) uses this pattern to illustrate how AI can shift where the thinking happens. As he explains, the LLM "reduces effort and provides an answer… relocating part of the thinking process outside the learner," which means the task gets done but less of the cognitive structure that supports that success is built.

Nosta points to two practices that help turn AI from shortcut to learning engine:

Iterative engagement
A back-and-forth interaction with the AI that builds understanding instead of bypassing it

Learner centricity:
Resources generated on the spot, tuned to a learner's interests and needs

[Good Things Foundation (2024)](#) emphasizes supported, reflective, personalized AI use in adult learning. All three perspectives converge on the same message: AI can improve performance, but only thoughtful design preserves learning.

What learning theory tells us

Learning theory helps connect what AI cannot do to why writing remains such an important cognitive activity.

Behaviourism

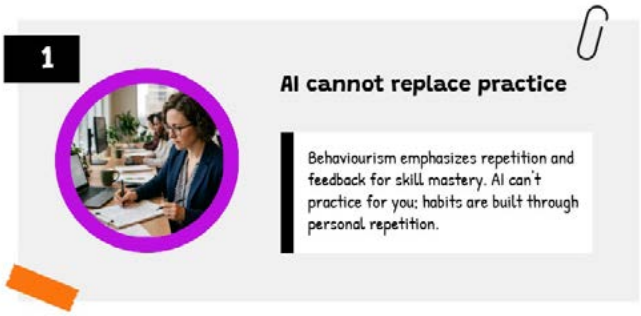

Behaviourism tells us that skills develop through repetition, reinforcement, and feedback. Mastery is reinforced when a learner performs the behaviour themselves.

AI can generate text, but it cannot do the practise for the learner. It cannot build habits or strengthen skill through comprehension on someone else's behalf.

Cognitivism

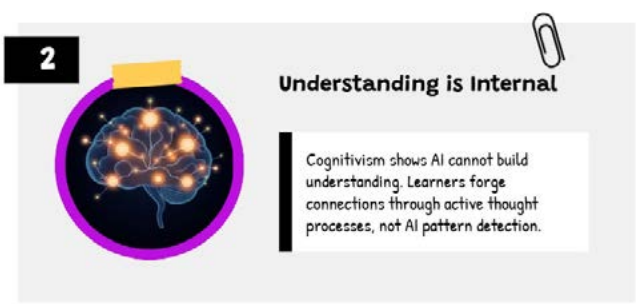

When learners write, they activate prior knowledge, organize ideas, and build schemas. The cognitive friction they experience such as pausing, sequencing, revising, and noticing gaps, is part of how understanding forms.

AI can produce concept-like responses by detecting patterns in data, but it cannot integrate new information with lived experience or build the learner's internal schema. It can support cognition by prompting reflection or offering explanations, but it cannot do the internal work of durable learning for the learner.

This matters because working memory is limited, and long-term memory becomes useful only when knowledge is stored and retrievable in real time. When learners outsource retrieval to AI, they may skip the consolidation that makes knowledge available under pressure.

If assessment is only measuring whether a learner can produce a correct answer, then AI makes that easier to fake. If assessment is measuring whether a learner has built the cognitive infrastructure to think with what they know, then the task is different. The answer is to design assessments that surface thinking, not just polished output.

Constructivism

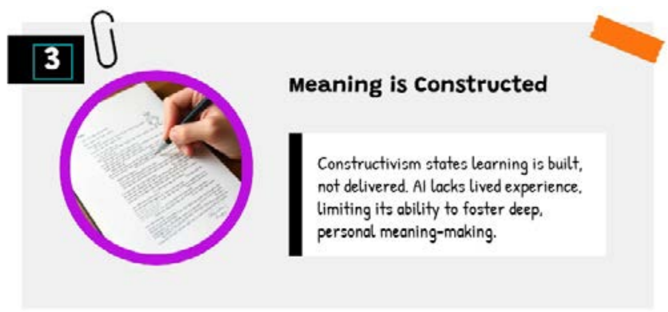

Constructivism reminds us that learning is not delivered but constructed. Learners build meaning through relevance, lived experience, dialogue, and community. Understanding grows as they connect new ideas to what they already know, test interpretations against their own lives, and negotiate meaning with others.

This is slow, relational, identity-shaping work, and it is where AI has clear limits. AI lacks lived experience, identity, perspective, and personal stake. It can imitate voice, but it cannot draw on memory, emotion, culture, or context in the way humans do when they make meaning.

Voice, perspective, and the ability to locate yourself in an argument are central tools for adults navigating work, relationships, and civic life. They are built through the act of choosing a stance, defending it, revising it, and understanding why it matters.

When AI generates a draft too early, it may shape the learner's stance before the learner has fully formed one. It can fill in the "I think…" before the learner has fully developed the thought. A learner who never had to defend an argument has missed a key moment where meaning is made.

The text may look polished, but the learner may not have built the internal architecture that supports reasoning, communication, and action in the world.

Edutopia emphasizes that authentic writing emerges from relationships, collaboration, and personal insight. Those are human processes, social, emotional, and contextual, that AI cannot fully automate.

AI can support meaning-making by offering prompts or examples, but it cannot replace the learner's construction of understanding. Without personal relevance and interpretation, meaning making is weakened.

Once we map learning theory onto AI literacy, the pattern is clear:

BEHAVIOURISM

➞ AI cannot replace the practice loops that build skill.

COGNITIVISM

➞ AI cannot replace the cognitive friction that builds understanding.

CONSTRUCTIVISM

➞ AI cannot replace the social and experiential processes that build meaning.

Why writing still matters

Learning theory helps clarify why writing still matters.

Behaviourism shows that writing builds skill through practice. Cognitivism shows that it deepens understanding through mental effort. Constructivism shows that it helps learners make meaning through personal and social engagement.

AI can generate text, but learners still have to practise, think, revise and interpret. Those are the steps that lead to lasting learning.

This is why writing remains essential. Writing is not just a product to be completed; it is a process through which learning happens.

AI provides speed, access and assistance. The learner provides effort, interpretation and meaning.

Learning depends on the learner, and that is something no technology can replace.

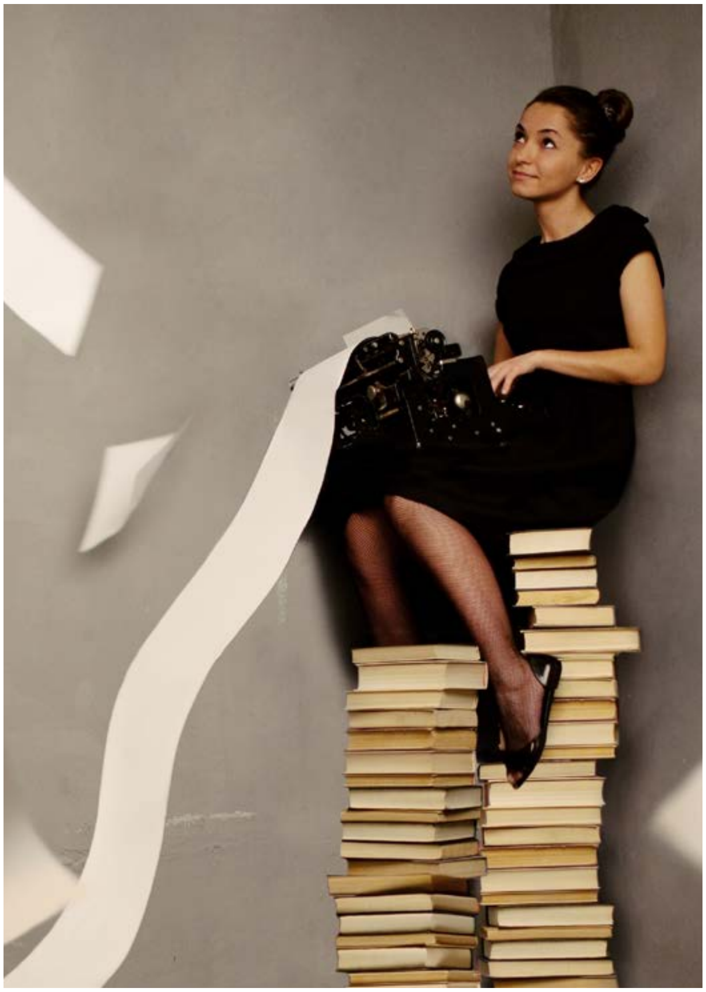

Which tool does what

AI tools serve different purposes. Some are best for drafting, some for research, some for embedded productivity and some for longer-form synthesis.

Which AI tool best fits your current workflow needs?

Screen capture this!

General-Purpose AI: Co-pilot, Gemini, Chat GPT	Perplexity AI*	Embedded AI in Microsoft 365 and Google Workplace	Claude
Drafting and brainstorming Ideal for creative and conversational tasks **Provides conversational output based on general training** Relies on user judgment for factual verification **Editing, explaining, and working through complex ideas** Suited for iterative content generation	**Functions as a research-first accelerator** Best for gathering evidence, trends, and statistics **Offers answers supported by source links** Enables users to confirm accuracy via references **Specialized for finding competitor info and hard-to-find background** Suited for evidence-based inquiries	Copilot supports you directly inside Microsoft 365 applications. Embedded in Outlook, Word, Excel, and PowerPoint Gemini supports Google Workplace Gemini supports you directly in Google Workplace. Embedded in Gmail, Docs, Sheets, and Slides Use for drafting, summarizing, generating slides, analyzing data, suggesting next steps	Operates independently of specific office suites Focuses on structured work and multi-step reasoning processes Handles large context and improves code quality Best for building tools, organizing documents, and drafting proposals

** Perplexity functions as a research accelerator, offering quick answers supported by source links. Although useful, it is fallible. The linked references help users confirm accuracy and treat its output as informed guidance rather than definitive authority.*

Mapping AI use cases, responsibilities and risks for learners

Different AI uses come with different expectations. Creative brainstorming can tolerate more flexibility, while research focused or automated tasks require closer verification. The graphic below highlights how increasing precision in the task also increases the need for careful validation.

- **General-purpose AI (low risk)**

  Use cases: *Clarify writing, draft responses, brainstorming partner.*

  These tasks live at the top of the graphic because they are generative and collaborative. The user looks for variations, tone adjustments or a starting point. The user is like a creative director.

- **Long-context tools (medium risk)**

  Use cases: *Pitch ideas, summarize threads, structure documents.*

  These tasks require the AI to ingest larger amounts of information and synthesize or format it. The risk is medium because if the AI misinterprets a long document, the user might misinform others. The user transitions from a creative director to an editor/fact-checker.

- **Research-first tools and embedded AI: Building simple tools (high risk)**

  Use cases: *Research support, build workflows.*

  When using AI for evidence gathering, the risk is high because LLMs may invent plausible-sounding facts, citations and sources. In a research-first context, the tool is a compass. The user must be the ultimate anchor.

When embedding AI into workflows to automate repetitive tasks, the risk is high due to data privacy, governance and the potential for compounding errors at scale. The graphic highlights some critical guardrails. The user is acting as a compliance officer.

AI use cases and responsibilities

Use Case	What AI Is For	User Responsibilities	Risk Level
Clarify writing	Simplify complex messages	Confirm meaning, tone, and intent remain accurate	Low
Draft responses	Generate first-draft options to accelerate writing	Edit for accuracy, appropriateness, and context	Low
Brainstorming partner	Explore options, variations, and creative directions	Select, refine, and own the final direction	Low
Pitch ideas	Produce structured drafts using best-practice patterns	Validate claims; ensure alignment with goals and audience	Medium
Summarize threads	Condense long email chains or documents	Cross-check summary against original content	Medium
Structure documents	Organize proposals, reports, and outlines	Review logic, accuracy, and completeness before publishing	Medium
Research support	Surface connections, trends, multilingual insights	Verify facts; check sources; cite appropriately	High
Build simple tools	Automate repetitive, low-risk tasks	Avoid processing private data; never automate decisions about people	High

New framework worth exploring

The [ASCD framework] emphasizes that an AI ready graduate is someone who can:

- **Learn with AI** but still set goals, seek feedback and stay accountable for their own growth
- **Research with AI** but evaluate claims, compare sources and detect contradictions
- **Synthesize with AI** but choose the right level, format and context for the audience
- **Solve problems with AI** but generate ideas, explore perspectives and break creative blocks
- **Connect with AI** but build human collaboration, bridge language gaps and coordinate teams
- **Tell stories with AI** but shape narratives, visuals and messages with intention and integrity

Across all six roles and the concepts and frameworks discussed in this bulletin, the through line is the same: **AI extends human capability, but it does not replace human responsibility.**

The most important question is not which tool is smartest, but which tool fits the task and the responsibility and risk level that come with it.

From theory to practice: Experimenting with AI intentionally

Google OPAL

Last, we promised GEMS and we're delivering. Gemini Gems are customized versions of Google Gemini that use specific instructions and knowledge files each time they generate a response. For example, an OALCF Find Information L3 expert Gem could include detailed information about the target audience, examples of the task, curriculum guidelines, and tone of the output, allowing learners to review texts and extract information without having to prompt the chatbot each time.

To close this bulletin, we're spotlighting a resource that shows what happens when AI stops acting like a black box and becomes something you can open up, inspect and actually teach with. This guide walks you through **Google Opal**, the engine behind a special kind of Gem.

What makes this resource worth your time is that it doesn't just show what Opal can do, it shows learners **how the system thinks**. Inside, you'll see:

- How a machine operates with workflows
- The instructions the AI is following
- How your inputs shape the outputs
- How learners can click **Editor** to watch the whole logic chain unfold

It's a practical, visual way to teach AI literacy, workflow thinking and responsible use of AI.

If you want a resource that helps learners move from "*AI gives me answers*" to "*I understand how this tool works*", download Exploring Google Opal now!

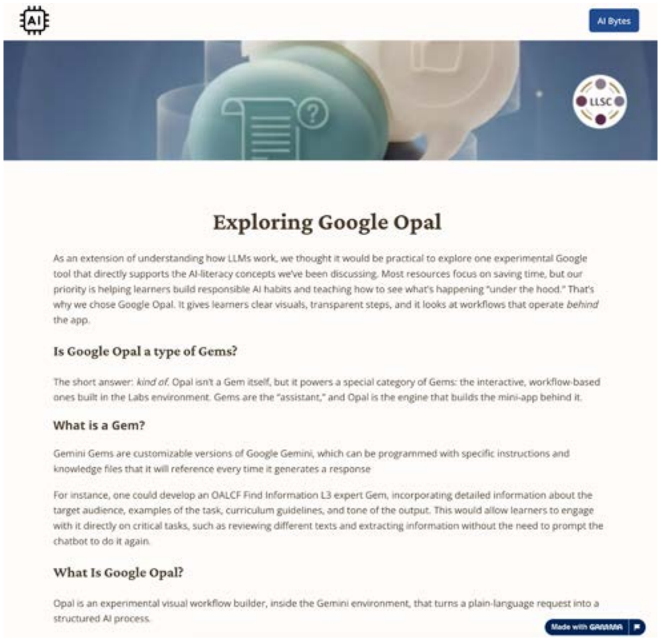

AI Bytes

LLSC

Exploring Google Opal

As an extension of understanding how LLMs work, we thought it would be practical to explore one experimental Google tool that directly supports the AI-literacy concepts we've been discussing. Most resources focus on saving time, but our priority is helping learners build responsible AI habits and teaching how to see what's happening "under the hood." That's why we chose Google Opal. It gives learners clear visuals, transparent steps, and it looks at workflows that operate *behind* the app.

Is Google Opal a type of Gems?

The short answer: *kind of.* Opal isn't a Gem itself, but it powers a special category of Gems: the interactive, workflow-based ones built in the Labs environment. Gems are the "assistant," and Opal is the engine that builds the mini-app behind it.

What is a Gem?

Gemini Gems are customizable versions of Google Gemini, which can be programmed with specific instructions and knowledge files that it will reference every time it generates a response

For instance, one could develop an OALCF Find Information L3 expert Gem, incorporating detailed information about the target audience, examples of the task, curriculum guidelines, and tone of the output. This would allow learners to engage with it directly on critical tasks, such as reviewing different texts and extracting information without the need to prompt the chatbot to do it again.

What Is Google Opal?

Opal is an experimental visual workflow builder, inside the Gemini environment, that turns a plain-language request into a structured AI process.

Made with GAMMA

In this edition, we dug into the mechanics of learning and how AI shapes and reshapes those associations. In our upcoming editions, we'll be shifting our focus to a deeper conversation raised by our community: AI and power. By power, we don't mean electricity, we mean control. We will explore who owns and builds the infrastructure behind these digital tools, whose specific worldviews their answers amplify, and whose political or economic interests these systems ultimately serve.

We hope this resource supported your teaching, and we can't wait to share what's coming next. Stay tuned for our upcoming podcasts!

References researched and attributions

Bastani, H., Bastani, O., Sungu, H., Ge, Ö., Kabakcı, R., & Mariman, R. (2024, June 25). *Generative AI without guardrails can harm learning: Evidence from high school mathematics*. PNAS, 122. https://doi.org/10.1073/pnas.2422633122

Celik, I., Zabolotna, K., & Viberg, O. (2026, February 16). Knowledge construction with GenAI: The role of theory-informed prompt engineering in achieving pedagogical alignment. *Innovations in Education and Teaching International*, 1–19. https://doi.org/10.1080/14703297.2026.2631586

Floridi, L., Morley, J., Novelli, C., and Watson, D. (2025, December 10). What kind of reasoning (if any) is an LLM actually doing? On the Stochastic Nature and Abductive Appearance of Large Language Models. *ArXiv*. https://arxiv.org/abs/2512.10080

Hudd, S. S., et al. (2025, March 7). Reading, writing, and thinking in the age of AI. *Faculty Focus*. https://www.facultyfocus.com/articles/teaching-with-technology-articles/reading-writing-and-thinking-in-the-age-of-ai/

Ritter, L. A. (2026, February). The impact of AI tools on adult learning outcomes in online education. *International Journal of Research and Review*, *Volume 13; issue 2* https://doi.org/10.52403/ijrr.20260213

Roberts, J. (2025, February 25). Authentic Writing in the Age of AI. Edutopia. https://www.edutopia.org/article/teaching-authentic-writing-age-ai/

Sedita, J. (2025, May 9). Writing instruction in the age of AI. *Keys to Literacy*. https://keystoliteracy.com/blog/writing-instruction-in-the-age-of-ai/

Stöffelbauer, A. (2023, October 24). *How Large Language Models Work: From Zero to ChatGPT*. Data Science at Microsoft, Medium. https://medium.com/data-science-at-microsoft/how-large-language-models-work-91c362f5b78f

University of Calgary Library. (2026, March 9). Artificial intelligence: AI literacy. *Libraries and Cultural Resources, University of Calgary*. https://libguides.ucalgary.ca/c.php?g=733971and p=5278508