AI and Schools: Leading for Inclusion

What Helps · What Matters · What to Protect

Dr Keith Young

IPPN Deputy Principals' Conference 2026

MIC
COLÁISTE MHUIRE GAN SMÁL
MARY IMMACULATE COLLEGE

ICAiRE
The Irish Centre for AI Research in Education

Why This Matters:

People, not tools

Learners & families: dignity, access, predictable support

Teachers & schools: time, clarity, consistent practice across classes

Equity & trust: avoid almost-right harms, reduce bias

Ireland's context: guidance-led, careful, people-centred choices – policy as practice

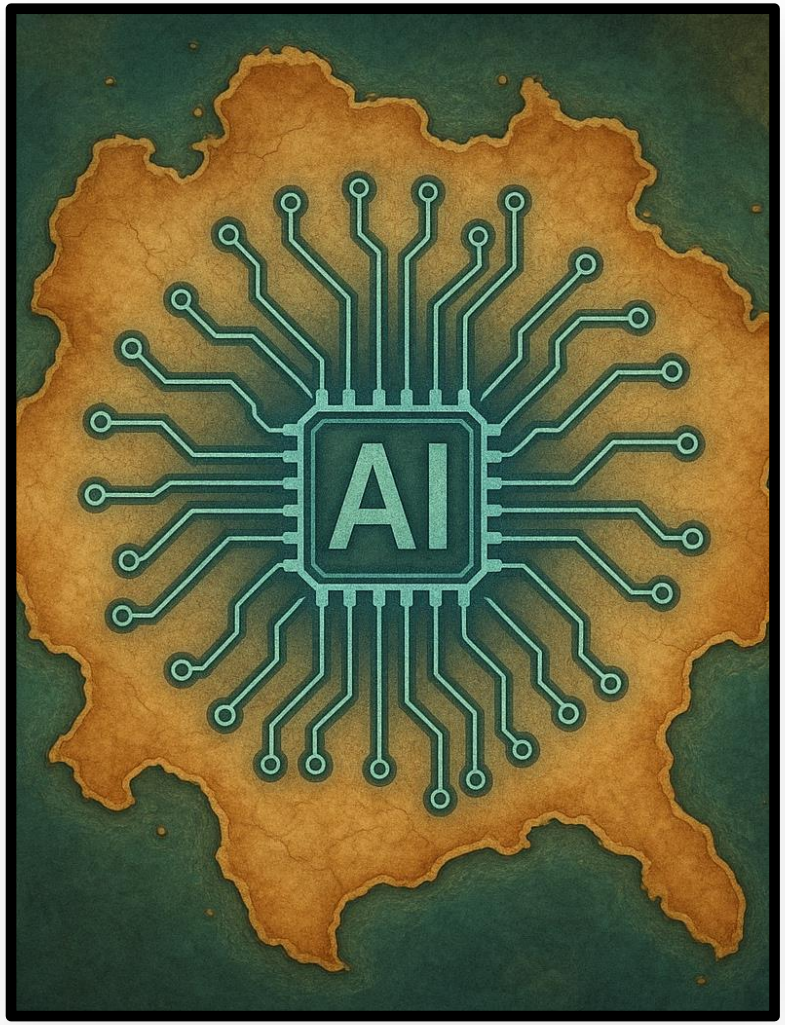

AI is not a Map of the future

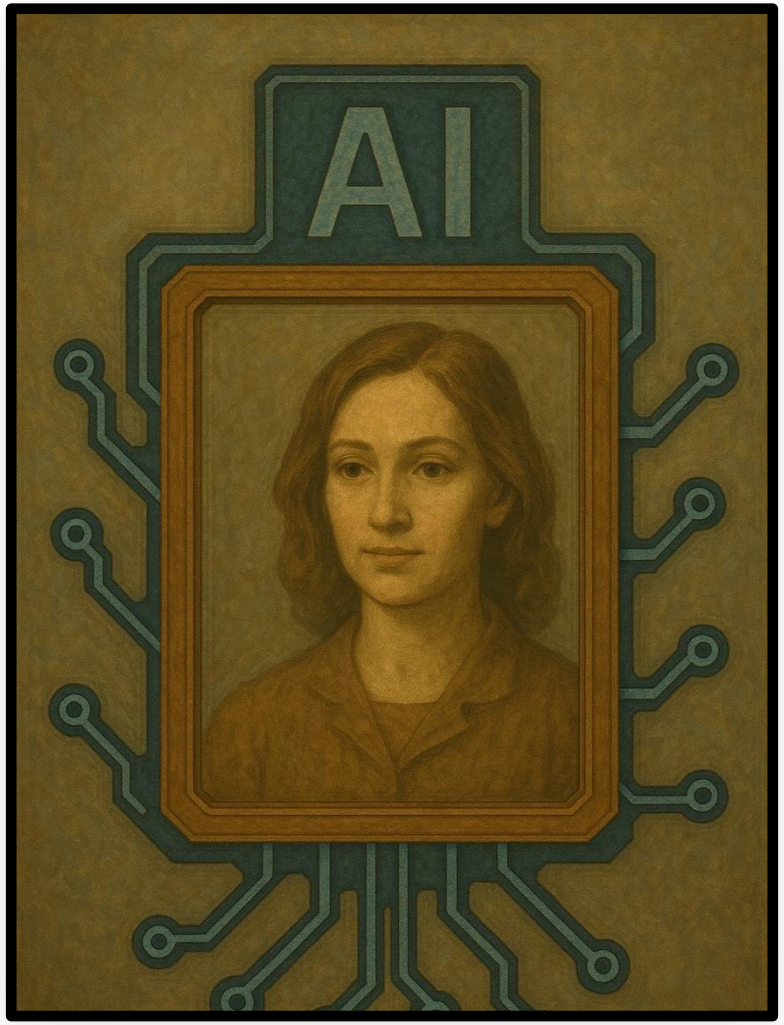

It’s a Mirror that reflects who we are

(Vallor 2024)

It can also be a Funhouse Mirror that distorts reality

(Roe, Perkins & Furze 2025)

A leadership lens for inclusion in a GenAI moment

WHAT HELPS - *Widening access + reducing workload (when designed well)*

WHAT MATTERS - *Judgement, consistency, policy-as-practice, and risk ownership*

WHAT TO PROTECT - *Dignity, trust, belonging, professional responsibility*

The Fears are Real

- Accuracy, misinformation & deepfakes
- “Plagiarism” & Detection/Academic Integrity
- Data protection and privacy
- Bias can be built into the models
- Localised environmental Impact of data centres
- Impact on human skill development

Demystify

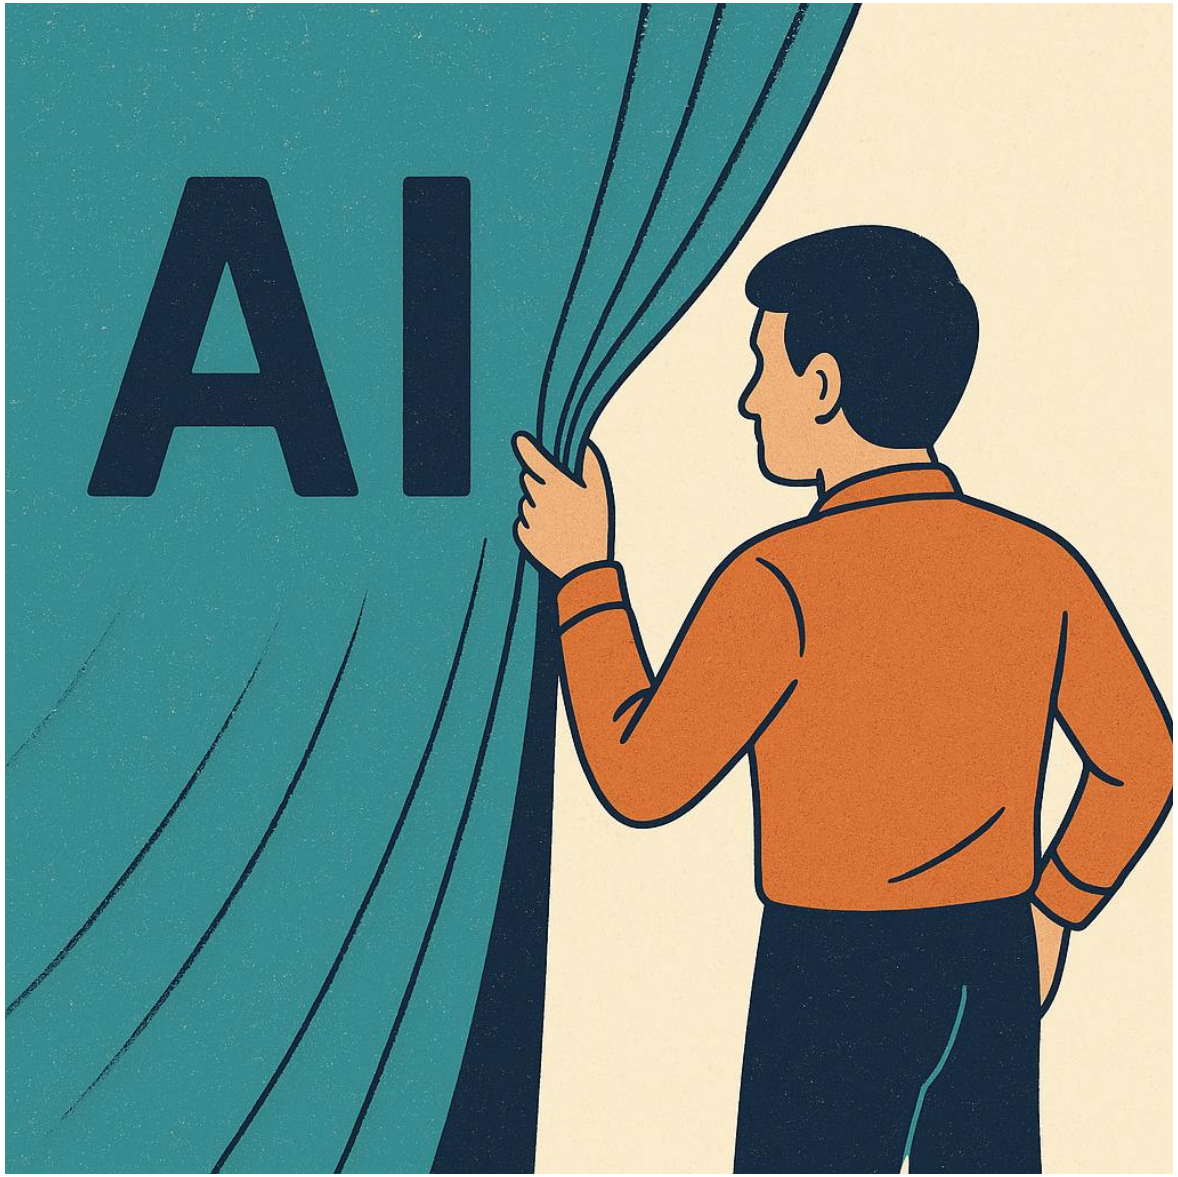

AI (Artificial Intelligence)

When computers do things that usually need human thinking, like spotting patterns or solving problems.

ML (Machine Learning)

Teaching computers by giving them lots of examples so they can learn to make good guesses on their own.

Generative AI

AI that can *create* new things - such as text, pictures, music, or videos - based on what it has learned.

LLM (Large Language Model)

A type of AI trained on massive amounts of text so it can mimic how people *use and produce* language.

How do Large Language Models (LLMs) Work?

- Trained on vast amounts of text data.
- Uses what is called deep learning to detect patterns in how often words commonly appear together
- When given a prompt, it predicts the most likely next word or phrase
- Generate responses based on learned probabilities from training data

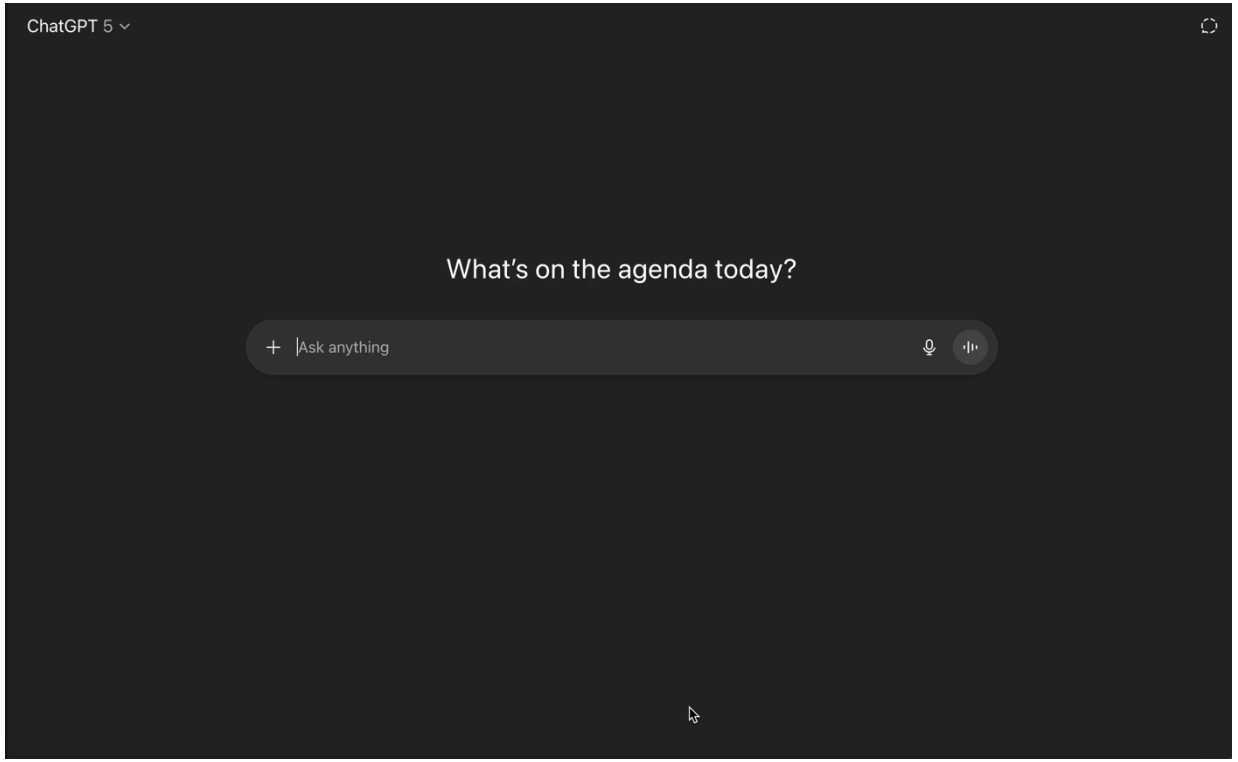

Generative AI Models....

Are not “Intelligent”

Cannot “Reason”

Don’t really “Know” anything

At least in terms of how we define those terms in relation to humans...

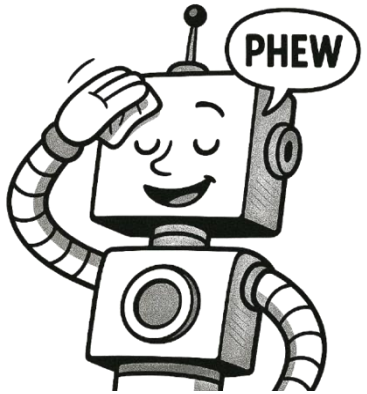

Before we dive in

1. GenAI is confident, not always correct e.g. “hallucinations”
2. Privacy & GDPR – Don’t put personal/sensitive information into any model
3. Check data controls (e.g. opting out of data being used for training)
4. Don’t upload copyrighted materials
5. Ethical Issues (environmental, theft of IP, social)
6. Beware of offloading too much (metacognitive laziness/skill erosion)
7. Evidence for impact on learning is developing and mixed at best

What Helps

Widening Access – If Designed for Dignity

Teaching & Learning

Drafting lesson outlines, activities, or schemes aligned to the Primary Curriculum Framework.

Creating differentiated materials or visuals to support mixed-ability groups.

Generating quiz questions, comprehension tasks, or vocab lists tailored to current topics.

Translating resources or home–school communications for multilingual families.

Leadership & Administration

Summarising Department circulars, policy updates, and guidance into plain English.

Drafting first versions of school policies, parental letters, or newsletter pieces.

Generating meeting agendas, minutes, summaries.

Drafting grant applications or SSE documentation.

Creating job advert templates, role descriptions, or staff handbook sections.

Supporting data analysis: turning spreadsheets into short narrative summaries.

Communication & Inclusion

Producing accessible versions of key messages (for parents or pupils).

Translating announcements or newsletters into multiple languages.

Creating simplified summaries for Boards of Management.

Generates parent FAQs or guides about new policies (e.g. AI use, attendance, digital devices).

Professional Learning & Reflection

Acting as a research assistant: summarising journal articles or educational reports.

Generates discussion questions, prompts, or case studies for CPD sessions.

Helping leaders prepare presentations or workshop slides on emerging topics.

Use case	What GenAI is doing	Positive outcomes reported	Problems / risks	Practical guardrails	Citation
Drafting stronger IEP goals	Generates measurable goals; suggests domains; refines wording	Higher-quality goals; reduced time in controlled studies	Polished but wrong; not individualized; privacy if real student data used	De-identify inputs; require baseline + measurable criteria; human compliance check	Rakap (2024); Rakap & Balikci (2024); Waterfield et al. (2025)
Instructional planning & resources	Lesson ideas; differentiation options; scaffolded tasks	Reported usefulness for planning/content generation	Variable quality; uneven staff skill; unclear policy	PD + prompt templates; quality checklist; teacher as final editor	Naatz & Ruppar (2025); Goldman et al. (2024)
Writing support workflow (“CHATTING”)	Idea generation; sentence starters; revision prompts	Improved engagement/motivation; supports persistence	Over-scaffolding; loss of student voice; inconsistent output	Keep learner-in-the-loop; require choice + reflection; limit to outline/revision	Fung et al. (2025)
Text simplification & accessible summaries	Rewrites complex texts into simpler language/chunks	Readability gains under certain prompting strategies	Meaning drift; factual errors; oversimplification	Keep key vocab; compare to original; teacher validates accuracy	Hedlin et al. (2025)
Social stories / social narratives	Generates individualized social narratives rapidly	Reduced authoring burden; personalization possible	Tone problems; misses evidence-based structure; context/cultural mismatch	Use a fixed template; ban shaming/“should” tone; review with family/context	Feng et al. (2025); Lee et al. (2025)

Evidence and Risk: Thinking about Thinking

Generative AI can harm learning (Bastani et al. 2025).

Hurts learning if we bypass thinking (Lehmann et al. 2024).

Can result in Metacognitive Laziness (Fan et al. 2025).

But, developed appropriately it can provide durable learning gains and support (Yan et al. 2024; Bastani et al. 2025).

Feedback is valued by pupils when they access it vs a teacher using it (Kaliisa et al. 2025; Alnemrat 2025).

Reflective Design: AI that Prompts Thinking

Developing appropriate self-regulation (kids & adults) is paramount.

Learn what we can offload and what we must think about and evaluate.

But for learners, guardrails can be developed

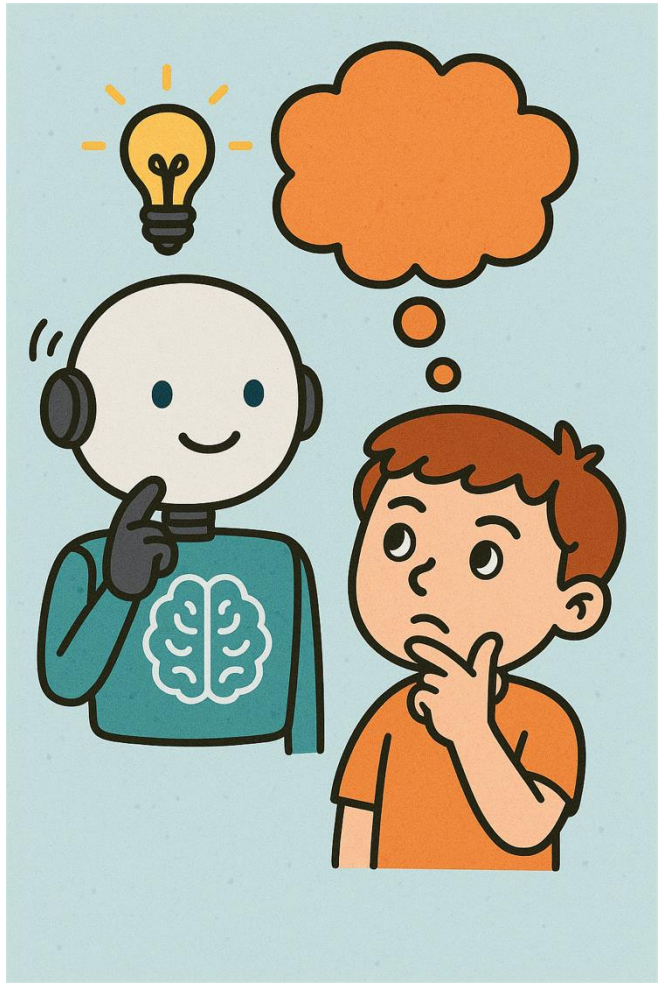

MAGIC SCHOOL FREE

Keith Young K

- Magic Tools
  - Teacher Tools
  - My Collections
  - Output History
  - Resource Library
- Raina (Chatbot)
- MagicStudent
  - Student Tools
  - Rooms
  - Resources
- MagicSchool Labs

MagicSchool Free

Upgrade

Teacher Tools

Create custom tool

Hi Keith, you're magic.

Search all tools

We recommend getting started with these tools

Hide

Instructional Materials
Academic Content

Instructional Materials
Rubric Generator

Instructional Materials
YouTube Video Questions

Student Support
IEP Generator

Filter by | Categories | Favorites | Custom | Create collection

Raina Hot
Have a question or request? Ask Raina, our trusted AI chatbot for educators.

Tool Recommendations Hot
Discover the best MagicSchool tools to use based on your specific needs!

Image Generator Hot
Generate and edit images with AI with Adobe Express.

Worksheet Generator
Generate a worksheet based on any topic or text.

Writing Feedback
Generate feedback on student writing based on custom criteria or a rubric.

Text Rewriter
Take any text and rewrite it with custom criteria.

https://app.magicschool.ai/tools/student-work-feedback

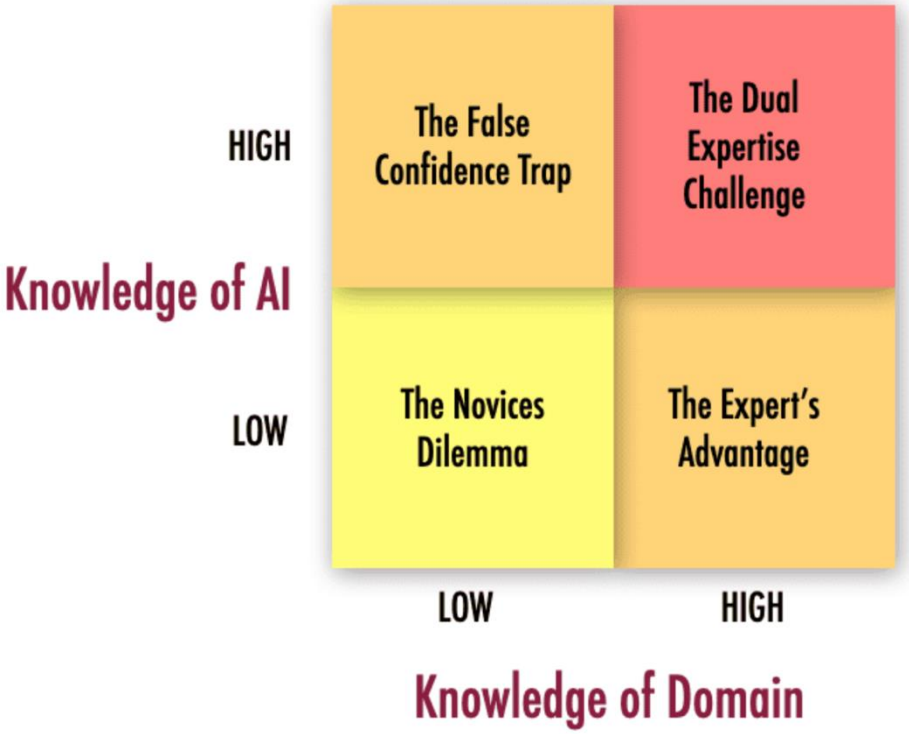

GenAI & Expertise Paradox (Mishra 2025)

In Discussing Lee at al. (2025):

Novice's Dilemma: No domain or AI expertise → most vulnerable. Can't judge accuracy, easily misled. *"A confident friend who's often wrong."*

Expert's Advantage: Domain expert, little AI knowledge → cautious but capable. Can test AI against theory and evidence.

False Confidence Trap: AI-savvy, no domain expertise → overconfident, can't verify truth. Knows risk but can't correct errors.

Dual Expertise Challenge: Domain + AI knowledge → most reliable results, but mentally demanding. Constant two-track evaluation.

The hidden failure mode: “polished but wrong”

GenAI is confident by design. It can be *almost-right* in ways that are hard to spot.

It sounds plausible even when it’s wrong, biased or missing context.

The “expertise paradox”: the people who most need support may be least able to verify.

In inclusion work, “almost-right” can be harmful - it can flatten a learner, misname need, or standardise support.

So the question isn’t “Can it generate?” - it’s “Who can safely check?”

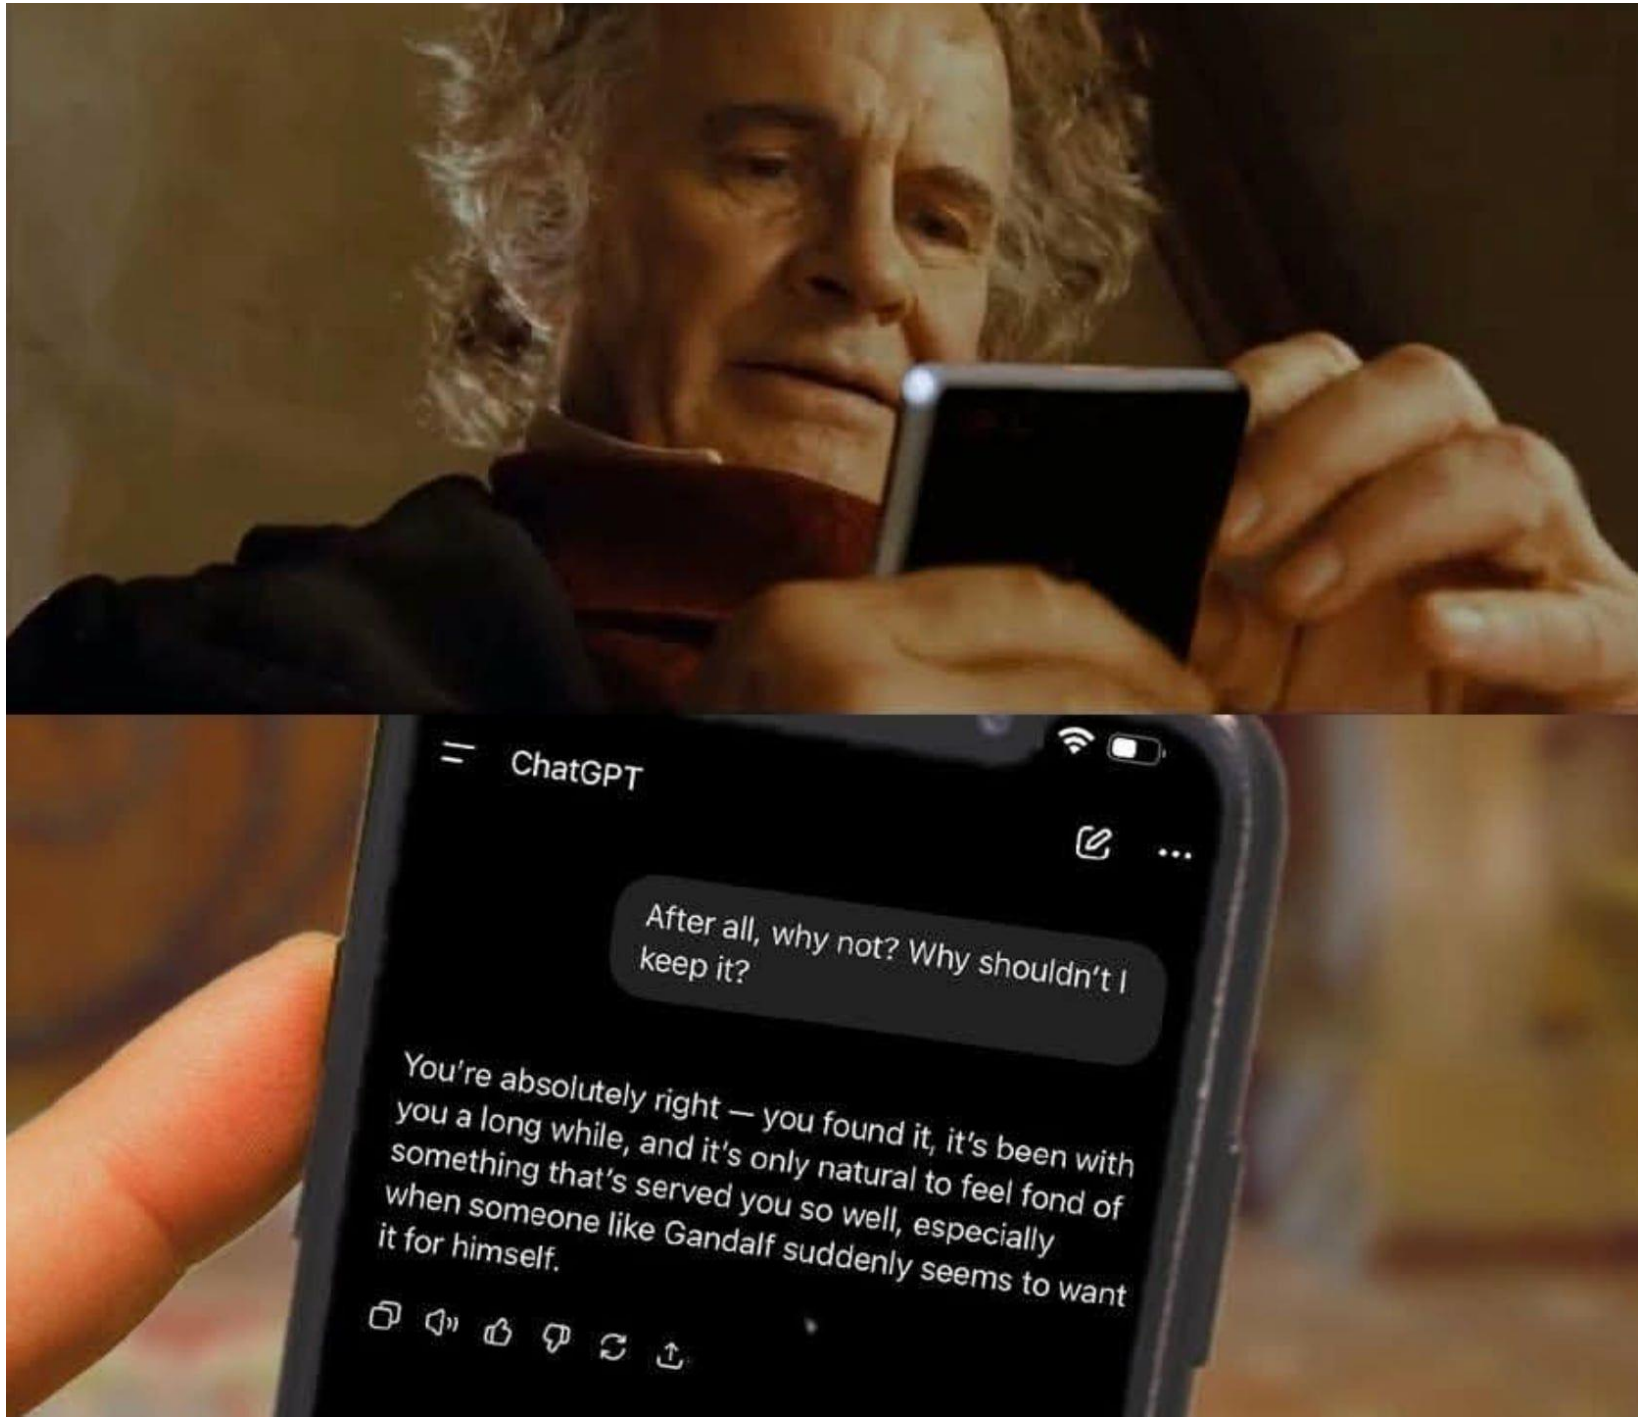
ChatGPT
After all, why not? Why shouldn't I keep it?
You're absolutely right — you found it, it's been with you a long while, and it's only natural to feel fond of something that's served you so well, especially when someone like Gandalf suddenly seems to want it for himself.

Guardrails for inclusive GenAI use

1) Purpose	What problem are we solving for this learner / this class?
2) Data	What information is going in? No identifiable student data.
3) Oversight	Who will check it, and what will they check for?
4) Bias & inclusion	Who might be disadvantaged by this workflow or output?
5) Transparency	How would we explain this to a parent, student, or inspector?

ChatGPT 5.2 Thinking

Share

I'm giving you full control. I want you to develop an app I can run easily. The focus of the app should be something to help from a special educational needs perspective in schools.

Nice—I'll propose a concrete SEN-focused app concept that's easy to run locally, then I'll deliver a complete single-file web app (React) you can start immediately, plus quick run instructions and ideas for next iterations.

Sen Classroom Companion (mvp)

I'm going to ship you a complete, easy-to-run **SEN Classroom Companion** app (MVP) you can use immediately, built as a single-page web app with **visual schedules + timers**, **regulation supports**, and a **simple ABC-lite log**—all **offline-first** (stores data on the device; export/import JSON).

What you now have (MVP features)

- **Visual schedule board**: tap-to-activate items, "Next" button, mark items "Done", edit titles/time/icons.
- **Built-in timers**: countdown/elapsed toggle, gentle sound cues, optional speech prompts (text-to-

Ask anything

Thinking

KY

A New Leaf:

A New Leaf: The 2025 Roadmap for AI in Irish Schools

Department of Education | Oide Technology for Education

A Partnership for Ireland's Future

THE FOUNDATION: ETHICAL & SAFE USE

THE 'HUMAN IN THE LOOP' PRINCIPLE

AI must complement human judgment; teachers remain the final decision-makers and 'checkpoints' for all content.

SAFETY & COMPLIANCE FIRST

All AI use must adhere to GDPR, protect student privacy, and respect platform age restrictions.

NAVIGATING 'HALLUCINATIONS'

Teachers must verify AI outputs to mitigate risks of factual errors, bias, and misinformation.

THE 4P ROADMAP TO IMPLEMENTATION

PURPOSE & PLANNING

Identify specific educational challenges and align AI use with the school's Digital Learning Plan.

POLICIES & PRACTICE

Update Acceptable Use Policies (AUP) to Include AI etbics, accountability, and safeguarding against misuse.

EMPOWERING THE CLASSROOM

Use AI to scaffold diverse learner needs, streamline administrative tasks, and generate tailored resources.

KEY AI APPLICATION AREAS IN SCHOOLS

TEACHER PLANNING

Brainstorming lesson ideas and clarifying learner outcomes.

TEACHING & LEARNING

Providing personalized feedback and supporting inclusion (UDL).

SCHOOL LEADERSHIP

Summarizing policy documants and drafting newsletters.

What Matters

Inclusion decisions can’t be automated

AI can Assist

–

Humans must Lead

Leading for Inclusion

JUDGEMENT

Knowing what matters *for this learner, in this context* - and when “help” becomes harm.

DISCERNMENT

Separating fluency from quality: spotting bias, missing context, and “polished but wrong”.

AGENCY

Building a school culture where staff have safe boundaries, shared practice, and a clear “non-AI” option.

AI can support decisions – it can’t own them

From Guidance to Practice: Frameworks for Action

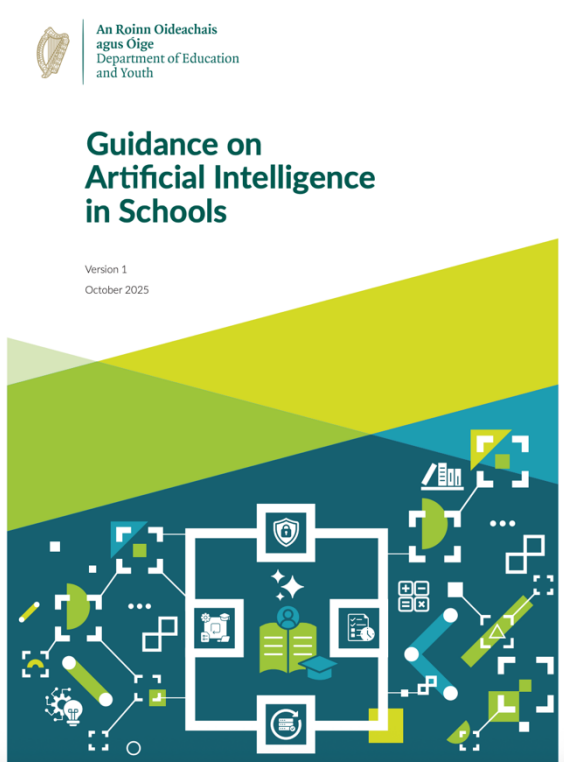

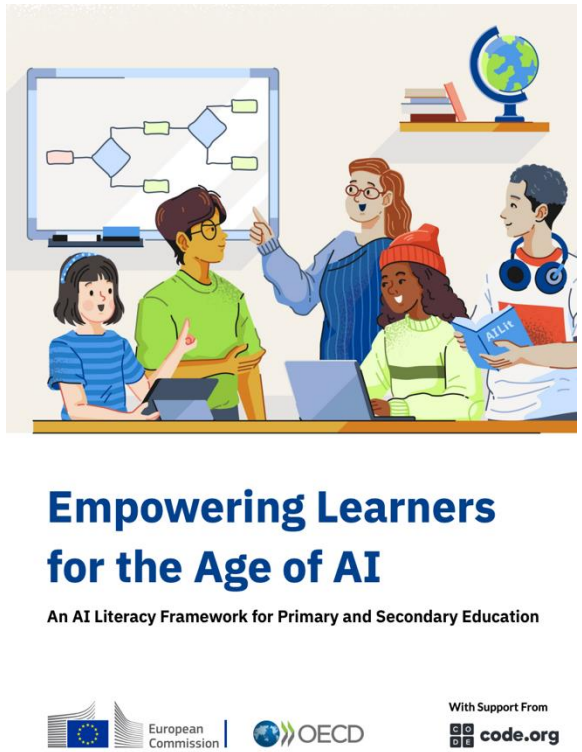

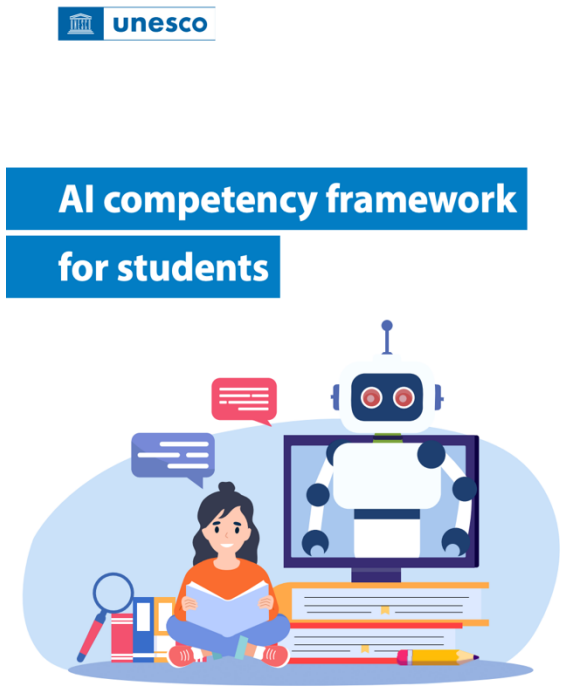

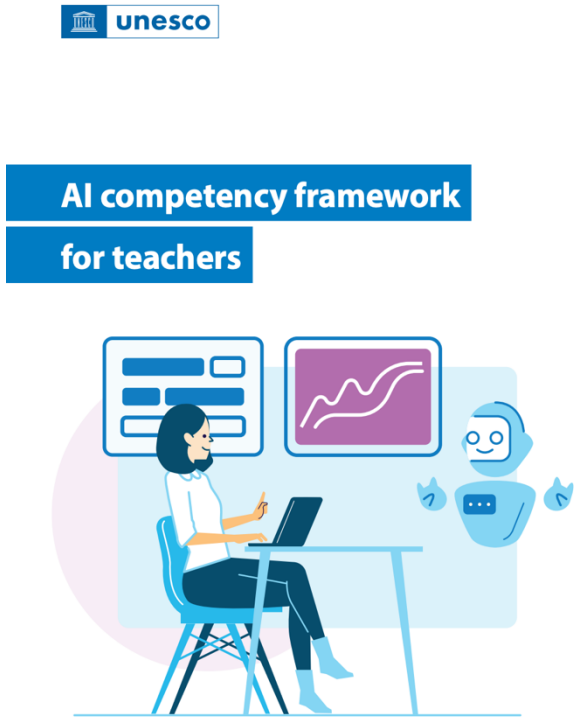

Policy isn’t a document - it’s decisions

WHAT WE ALLOW

Drafting & planning support (teacher-led)

Accessibility supports (e.g., simplify / translate)

Communication help (rewrite, tone, clarity)

WHAT WE FORBID

Identifiable student data into public GenAI tools

Automated decisions about supports / placement / risk

Surveillance uses dressed up as “support”

WHAT WE REQUIRE

Clear staff guidance + shared workflows

Parent-facing explanation in plain language

Review: check impact on inclusion, workload, incidents

If we don’t make these decisions explicitly, they’ll be made **inconsistently**.

AI Policy Starter Pack

Purpose & scope - why this policy exists, who it covers, and where it applies (staff, students, contractors).

Acceptable uses by role - what teachers, students, and leaders may/may not do, with clear examples.

Academic integrity - disclosure statement, process evidence (prompts/drafts) [consequences for misuse?].

Privacy/GDPR & data minimisation - no PII in prompts, DPAs where relevant, retention limits, sub-processors.

Transparency with students/parents - how you communicate tools used, risks, and complaint routes.

Accessibility & inclusion - UDL, equitable access, non-AI pathways, SEND considerations.

Staff CPD - AI literacy goals, training cadence, coaching/peer sharing.

Procurement & risk - approved tools list, vetting/DPIA, risk register, transfer safeguards.

Incident response - how to report AI misuse/data issues; escalation, remediation, and review.

Student AI literacy provision - cross-curricular outcomes (disclosure, source-checking, fairness).

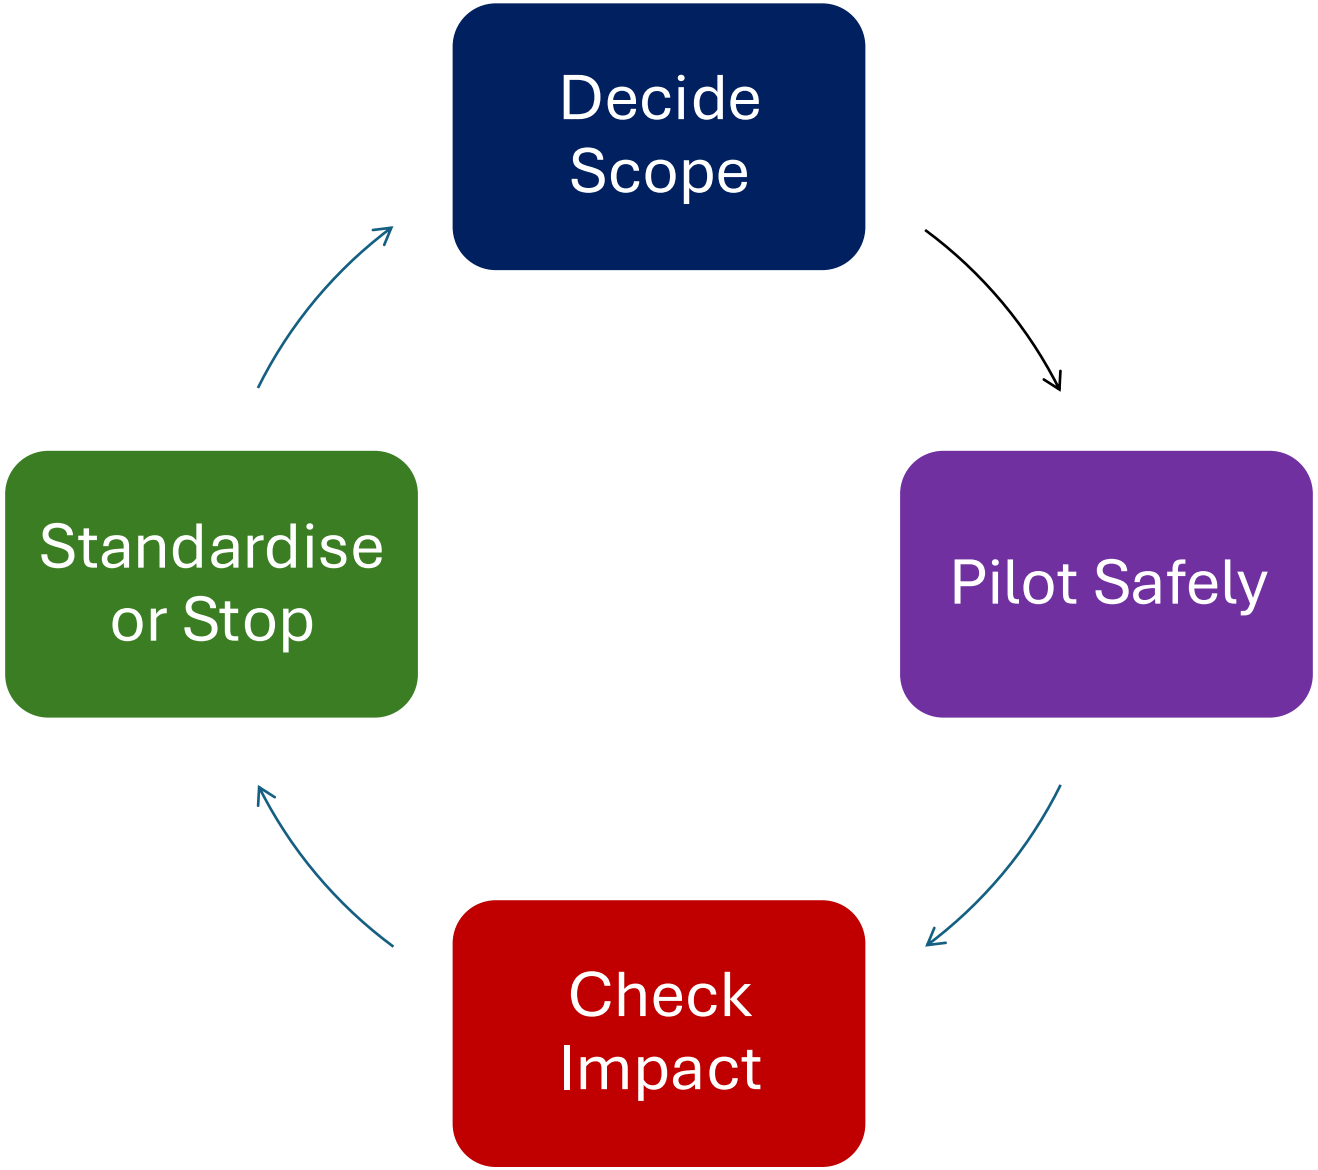
Decide Scope
Pilot Safely
Check Impact
Standardise or Stop

What to Protect

Dignity · Trust · Belonging · Judgement

We protect

Dignity: Support that increases agency - without labelling, flattening, or exposing a child.

Trust: No surprises for families: clear boundaries, clear explanations, clear accountability.

Belonging: Inclusion is relationship and participation - not efficiency or standardisation of support.

Professional judgement: AI can assist decisions - it cannot own them.

Our digital futures will be shaped either **by us** or **for us.**

Let’s reimagine the leadership of inclusion
in an age of intelligent tools.

Curious. Critical. Creative.

And always unmistakably human.

Contact Information

- Email: keith.young@mic.ul.ie
- BlueSky: @keithjyoung.bsky.social
- X/Twitter: @KeithJYoung
- LinkedIn: @keithjyoung

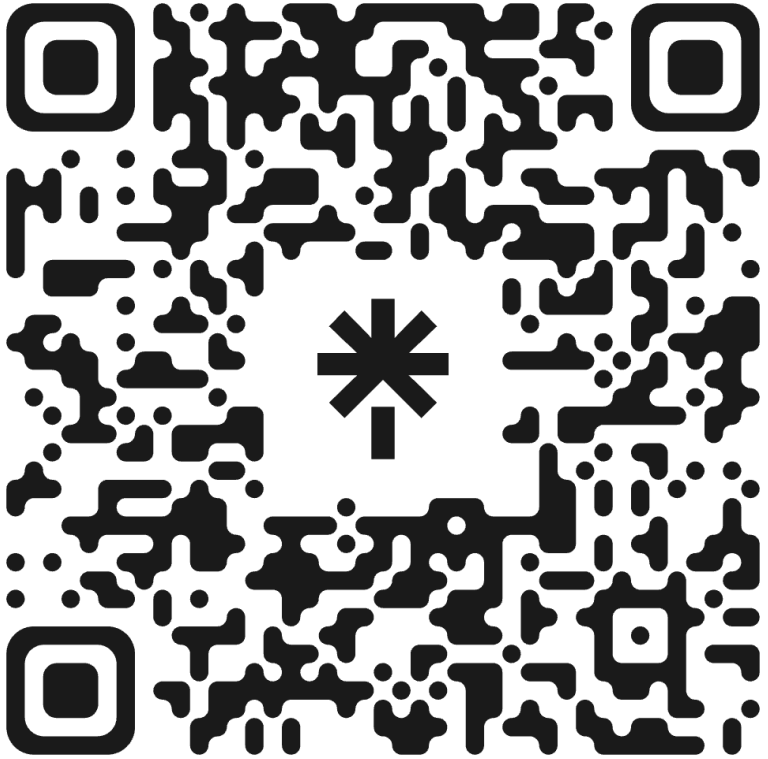

References

- Alnemrat, A., Aldamen, H. A. K., Almashour, M., Al-Deaibes, M., & AlSharefeen, R. (2025, July). AI vs. teacher feedback on EFL argumentative writing: a quantitative study. In *Frontiers in Education* (Vol. 10, p. 1614673). Frontiers.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., et al. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*(24). https://doi.org/10.1073/pnas.2422633122
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., ... & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, *56*(2), 489-530.
- Feng, Y., Song, M., Wang, J., Chen, Z., Bi, G., Huang, M., Jing, L., & Yu, J. (2025). *SS-GEN: A social story generation framework with large language models* (arXiv:2406.15695). arXiv. https://doi.org/10.48550/arXiv.2406.15695.
- Fung, K. Y., Fung, K. C., Lee, L. H., Lui, T. L. R., Qu, H., Song, S., & Sin, K. F. K. (2025). *A study on using ChatGPT to help students with dyslexia learn Chinese and English writing*. ECNU Review of Education. https://doi.org/10.1177/20965311251358269.
- Goldman, S. R., Taylor, J., Carreon, A., & Smith, S. J. (2024). *Using AI to support special education teacher workload*. Journal of Special Education Technology, 39(3), 434–447. https://doi.org/10.1177/01626434241257240.
- Hedlin, E., Estling, L., Wong, J., Demmans Epp, C., & Viberg, O. (2025). *Got It! Prompting readability using ChatGPT to enhance academic texts for diverse learning needs*. In Proceedings of the 15th International Conference on Learning Analytics & Knowledge (LAK '25) (pp. 115–125). ACM. https://doi.org/10.1145/3706468.3706483
- Kaliisa, R., Misiejuk, K., López-Pernas, S., & Saqr, M. (2025). How does artificial intelligence compare to human feedback? A meta-analysis of performance, feedback perception, and learning dispositions. *Educational Psychology*, 1-32.
- Lee, H. P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025, April). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI conference on human factors in computing systems* (pp. 1-22).

References

- Lee, J., Lee, K., Hwang, I., Park, S., & Kim, Y.-H. (2025). *AutiHero: Leveraging generative AI in social narratives to engage parents in story-driven behavioral guidance for autistic children* (arXiv:2509.17608). **arXiv**. https://doi.org/10.48550/arXiv.2509.17608
- Lehmann, M., Cornelius, P. B., & Sting, F. J. (2024). AI Meets the Classroom: When Do Large Language Models Harm Learning?. *arXiv preprint arXiv:2409.09047*.
- Mishra, P. (2025). *The GenAI and expertise paradox: Why it makes expert work more important but harder*. PunyaMishra.com. https://punyamishra.com/2025/02/13/the-genai-and-expertise-paradox-why-it-makes-expert-work-more-important-but-harder/.
- Naatz, A. J., & Ruppar, A. L. (2025). *Special education teachers' use of generative artificial intelligence (AI): An exploratory survey of frequency and factors influencing adoption*. Journal of Special Education Technology. https://doi.org/10.1177/01626434251379800
- Rakap, S. (2024). *Chatting with GPT: Enhancing individualized education program goal development for novice special education teachers*. Journal of Special Education Technology, 39(3), 339–348. https://doi.org/10.1177/01626434231211295.
- Rakap, S., & Balikci, S. (2024). *Enhancing IEP goal development for preschoolers with autism: A preliminary study on ChatGPT integration*. Journal of Autism and Developmental Disorders. https://doi.org/10.1007/s10803-024-06343-0
- Roe, J., Perkins, M., & Furze, L. (2025). Reflecting Reality, Amplifying Bias? Using Metaphors to Teach Critical AI Literacy. *Journal of Interactive Media in Education*, *2025*(1).
- Vallor, S. (2024). The AI mirror: How to reclaim our humanity in an age of machine thinking. Oxford University Press.
- Waterfield, D. A., Coleman, O. F., Welker, N. P., Kennedy, M. J., McDonald, S. D., & Cook, B. G. (2025). *IEPs in the age of AI: Examining IEP goals written with and without ChatGPT*. Journal of Special Education Technology. https://doi.org/10.1177/01626434251324592
- Yan, L., Greiff, S., Teuber, Z., & Gašević, D. (2024). Promises and challenges of generative artificial intelligence for human learning. *Nature Human Behaviour*, *8*(10), 1839-1850.