**RESEARCH BRIEF FOR EDUCATORS**

# When AI Supports Learning and When It Replaces Thinking

*An AI research brief for educators*

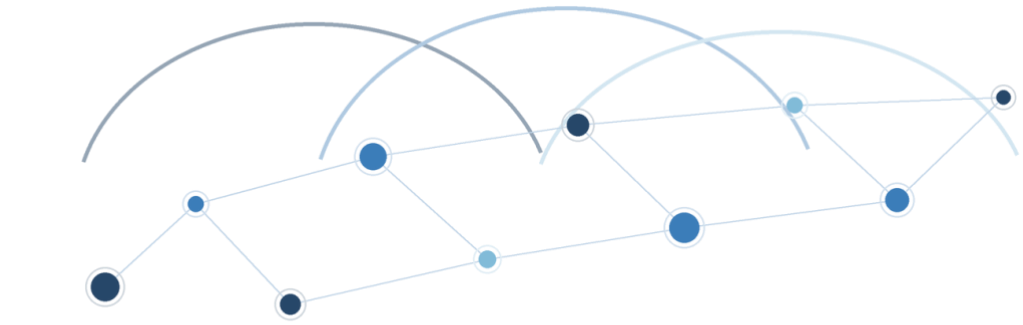

**By Med Kharbach, PhD**

www.medkharbach.com

# Introduction

> **What is the student still doing cognitively when AI is involved in the task?**
>
> A completed assignment tells us very little on its own. A student can produce a stronger essay, solve more practice problems, or finish an activity faster with AI. Yet the polished result may hide what happened during the learning process.

A year and a half ago, I started sharing reviews of research on AI in education at medkharbach.com/blog. My purpose was simple: to help fellow teachers, educators, and researchers keep up with the growing body of research on AI and what it means for our work. In each review, I try to move beyond summarizing the paper and draw out practical insights that we can use in our classrooms, research, and professional practice.

Since then, the AI Research Hub has grown immensely. I have now reviewed more than 200 papers, and as I worked through this expanding body of research, certain questions and themes began to appear repeatedly. Some deserved to be brought together and examined across several studies instead of within a single paper.

This is what led me to create the Med Kharbach AI Research Brief, which I will share periodically with subscribers to my newsletter. Each brief will focus on an issue relevant to educators and researchers, bring together what current research tells us about it, and translate the findings into practical ideas. I will also include activities, reflection questions, frameworks, and other resources that you can adapt to your own context.

For this first AI Research Brief, I am looking at the impact of AI on learning, particularly what happens to students' thinking when they use generative AI. The question guiding the brief is one I keep returning to in my own reading: **When does AI support learning, and when does it begin to replace the thinking students need to do for themselves?**

The research papers I am drawing on for this AI Research Brief are:

- *Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics* (2025), by Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman.
- *AI Assistance Reduces Persistence and Hurts Independent Performance* (2026), by Grace Liu, Brian Christian, Tsvetomira Dumbalska, Michiel A. Bakker, and Rachit Dubey.

- *ChatGPT as a Cognitive Crutch: Evidence from a Randomized Controlled Trial on Knowledge Retention* (2025), by André Barcaui.
- *Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance* (2025), by Yizhou Fan, Li Tang, Huyen Le, Kexin Shen, Shufang Tan, Yueying Zhao, Yuan Shen, Xinyu Li, and Dragan Gašević.
- *The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers* (2025), by Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson.
- *Summarise, Elaborate, Try Again: Exploring the Effect of Feedback Literacy on AI-Enhanced Essay Writing* (2025), by Brittany Hawkins, David Taylor-Griffiths, and Jason M. Lodge.
- *Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task* (2025), by Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes.
- *AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking* (2025), by Michael Gerlich.
- *Artificial Intelligence, Cognitive Offloading and Implications for Education* (2026), by Jason M. Lodge and Leslie Loble.

## The Central Finding

**AI outcomes depend on how assistance is designed.**

AI supports learning when it helps students engage more deeply with the intellectual work of a task. It begins to replace thinking when it removes the effort, decisions, and monitoring through which learning develops.

## What the research is showing

The evidence does not support a simple conclusion that AI is either beneficial or harmful. Outcomes depend on the task, the design of the tool, the instructions students receive, and what happens after AI assistance is removed.

**FINDING 01**

**A better product can conceal weaker learning**

Several studies found that AI improved the work produced while assistance was available. Those gains did not consistently carry into knowledge, transfer, retention, or independent performance (Bastani et al., 2025; Barcaui, 2025; Fan et al., 2025; Liu et al., 2026).

**FINDING 02**

**Quick answers may weaken persistence**

Brief access to complete AI answers was followed by lower unassisted performance in randomized experiments. Participants were also more likely to stop or skip in two of the three experiments (Liu et al., 2026).

**FINDING 03**

**The design of support changes the outcome**

Guardrails, hints, questions, feedback, and independent reconstruction keep more of the learning work with the student. Unrestricted answer generation makes cognitive offloading easier (Bastani et al., 2025; Fan et al., 2025; Hawkins et al., 2025).

**FINDING 04**

**Judgment is now part of the learning task**

Productive use requires students to decide what help they need, evaluate responses, verify claims, reject weak suggestions, and integrate useful feedback (Lee et al., 2025; Hawkins et al., 2025).

## STUDY 01: A Better Product Can Conceal Weaker Learning

### Bastani and colleagues: high school mathematics

Bastani and colleagues studied nearly 1,000 students in Grades 9, 10, and 11 at a high school in Türkiye. Classes were assigned to regular instruction without an AI tutor, GPT Base, or GPT Tutor. GPT Base resembled a standard ChatGPT interface. GPT Tutor included prompts and safeguards intended to support learning.

Students completed assisted practice and then took a closed-book, closed-laptop examination. During practice, both AI groups performed considerably better than the control group. Grades increased by 48% in the GPT Base condition and 127% in the GPT Tutor condition.

> **WHAT HAPPENED WITHOUT AI**
>
> The result changed when students worked independently.
>
> Students who had used GPT Base scored 17% lower on the unassisted examination than students in the control condition. Conversation logs showed that students frequently requested answers and copied solutions.

The AI helped them complete the practice problems, but it also allowed them to bypass some of the reasoning those problems were designed to develop. GPT Tutor produced a more encouraging result: its safeguards largely eliminated the negative examination effect, although it did not create a significant learning advantage over the control condition.

### What educators can take from it

- Assisted performance shows what a student can produce while AI is available.
- Independent performance shows what the student can do once that assistance is removed.
- If our goal is learning, we need evidence from both conditions.

*The experiment measured short-term outcomes in one subject and one school. It does not tell us what every AI tutor will do in every classroom. It does offer unusually strong causal evidence that the design of AI assistance matters.*

# STUDY 02: Quick Answers May Reduce Persistence

**Liu and Colleagues: Mathematics and Reading**

Liu and colleagues conducted three randomized experiments involving 1,222 participants. The tasks included fraction problems and reading-comprehension questions. Participants received brief access to GPT-5 during practice and then completed additional tasks without it.

AI assistance improved performance while it was available. Once it was removed, participants who had received that assistance generally performed worse and were more likely to skip questions.

> **The unassisted performance gap appeared across tasks.**
>
> In the first mathematics experiment, the AI group solved 57% of the later problems, compared with 73% in the control group. In the reading experiment, the corresponding figures were 76% and 89%. A larger mathematics replication found a smaller gap: 71% and 77%.

Skipping also increased significantly in the first mathematics experiment and in the reading experiment. The overall skipping difference was not statistically significant in the larger replication.

An exploratory analysis in the second experiment found that participants who requested direct answers later performed worse than participants who used AI for hints or clarification. This comparison was observational within the AI condition, so it cannot establish that answer requests caused the difference.

## What educators can take from it

When AI responds immediately and completely, students lose an opportunity to remain with the problem. Turning outward at the first sign of difficulty may weaken the persistence through which skills develop.

*This paper is a preprint under review. It measured immediate performance following brief AI exposure. Claims about lasting changes in motivation or learning would go beyond the evidence.*

# STUDIES 03–04: Ease During Study Can Carry a Learning Cost

## Barcaui: Knowledge Retention After 45 Days

Barcaui randomly assigned 120 Brazilian undergraduate business students to study material about artificial intelligence with either ChatGPT or traditional, non-AI resources. Approximately 45 days later, they received an unexpected retention test.

**Students in the traditional group remembered more.**

Among the 85 students who completed the delayed assessment, the traditional-learning group answered 68.5% of the questions correctly. The ChatGPT-assisted group answered 57.5% correctly. The difference was statistically significant, with a moderate effect size.

*Nearly 30% of the original participants did not complete the delayed test, although attrition was similar across conditions. The study involved one university, one field, one learning activity, and one version of ChatGPT. It measured retention without assessing possible benefits such as synthesis, efficiency, collaboration, or prompt development.*

## Fan and colleagues: Improved Essays, Limited Transfer

Fan and colleagues randomly assigned 117 university students to complete an English reading-and-writing task with ChatGPT-4, a human writing expert, a writing-analytics checklist, or no additional support. The ChatGPT group achieved significantly greater essay-score improvements than the other groups.

The groups did not differ significantly in knowledge gain, knowledge transfer, or post-task intrinsic motivation. The researchers also observed different patterns of planning, monitoring, and evaluation and raised the possibility of metacognitive laziness.

*The study did not directly measure metacognitive laziness with a mature instrument. Its single task, modest sample, and short duration also limit generalization.*

# STUDIES 05–06: AI changes where critical thinking happens

## Lee and Colleagues: Knowledge Work

Lee and colleagues surveyed 319 knowledge workers who regularly used generative AI. Participants described 936 examples from their work. Higher confidence in AI was associated with less reported critical-thinking effort, while greater confidence in one's own ability was associated with more critical thinking.

Participants described shifts from gathering information to verifying it, from solving a problem to integrating an AI response, and from completing a task to supervising the task and its output.

*The study was based on self-reports and associations. It cannot show that AI caused a reduction in critical thinking. Participants may also have confused reduced effort in general with reduced critical-thinking effort.*

## Hawkins and Colleagues: Feedback Literacy

Hawkins, Taylor-Griffiths, and Lodge observed 32 psychology students during a 25-minute simulated essay task. Students used AI to generate ideas, clarify the task, request feedback, paraphrase, evaluate suggestions, refine prompts, and sometimes deliberately stop using AI.

> **Productive use depended on student judgment.**
>
> All but one participant evaluated an AI response at least once, yet fewer than one-third cross-checked information against another source. Half deliberately avoided AI at some point in the task.

Self-reported feedback literacy was the only significant positive predictor of essay grade. The overall regression model was not statistically significant, and the sample was small, so the authors appropriately describe the study as exploratory.

# STUDIES 07-08: Cognitive Offloading Needs Careful Interpretation

## Kosmyna and Colleagues: Preliminary Neural Evidence

The preliminary “Your Brain on ChatGPT” study followed 54 participants across three essay-writing sessions; 18 completed a fourth session in which conditions changed. The researchers combined EEG, linguistic analysis, interviews, and essay scoring.

The Brain-only group showed the strongest brain-network connectivity, the search group was intermediate, and the LLM group showed the weakest. Participants in the LLM group also reported lower ownership and had more difficulty quoting their essays shortly after writing.

## Gerlich: Associations in a UK Sample

Gerlich surveyed 666 participants and interviewed a subset of 50. Greater reported AI use was associated with more cognitive offloading and lower critical-thinking scores. Because the design is correlational, it cannot establish whether AI use reduces critical thinking, whether people with different critical-thinking habits use AI differently, or whether other factors explain the relationship.

## Lodge and Loble: A Synthesis for Education

Lodge and Loble’s 2026 report brings this research together around a useful distinction. Offloading can free attention for other work, but detrimental offloading occurs when students outsource the mental activity needed to build knowledge, skills, and judgment. The report emphasizes purposeful pedagogy, well-designed tools, and teacher expertise.

# INTERPRETATION

## The Boundary Educators Need to Watch

| AI SUPPORTS THINKING | AI REPLACES THINKING |
|---|---|
| ✓ Follows an initial student attempt | × Supplies complete answers too early |
| ✓ Provides hints and graduated support | × Performs the target reasoning |
| ✓ Invites comparison and evaluation | × Encourages acceptance without checking |
| ✓ Returns responsibility to the learner | × Removes opportunities to persist |
| ✓ Is followed by independent transfer | × Produces work never tested independently |

## What counts as productive cognitive offloading?

Cognitive offloading is part of ordinary intellectual life. We use calculators, notes, diagrams, reference books, calendars, and search engines. The educational question is what gets offloaded and what the learner does with the capacity that becomes available.

### Productive examples

- Checking spelling after the student has composed an argument.
- Generating another example after the student explains the concept.
- Receiving a hint after making an initial attempt.
- Obtaining feedback that the student must evaluate and apply.

### Riskier examples

- Generating the first interpretation before the student reads closely.
- Supplying a complete solution before the student attempts the problem.
- Summarizing material the student is expected to understand and remember.

# THINK • USE • CHECK • RETURN

**A practical framework for classroom use**

**01**

**Think: establish a cognitive starting point**

Students attempt the problem, retrieve what they remember, make a prediction, develop a preliminary outline, or identify where they are uncertain before AI shapes the work.

**02**

**Use: define the form of assistance**

Ask AI for a hint, guiding question, error explanation, feedback against criteria, contrasting example, or challenge to an assumption. The requested role should match the learning goal.

**03**

**Check: make evaluation visible**

Students identify what is supported, what is missing, what conflicts with course material, what needs verification, and which suggestions they will accept, reject, or revise.

**04**

**Return: require independent reconstruction**

Students solve a related problem, explain the concept in their own words, reconstruct the argument from memory, transfer it to a new situation, or defend a decision without AI.

**DESIGN PRINCIPLE**

**Protect the thinking that the learning goal depends on.**

AI can remove peripheral friction. Students still need to perform the retrieval, interpretation, reasoning, evaluation, or judgment the lesson is intended to develop.

# BEFORE ASSIGNING AI-SUPPORTED WORK

## Educator checklist

Use this checklist while planning an activity or reviewing an existing assignment.

### The learning goal

- ☐ I can name the knowledge, skill, or judgment students should develop.
- ☐ I have identified the cognitive processes essential to that goal.
- ☐ I know what students should eventually be able to do independently.

### The role of AI

- ☐ The tool has a defined role: hint, question, feedback, example, or explanation.
- ☐ AI does not perform the central intellectual work of the task too early.
- ☐ The level of assistance matches students' prior knowledge and readiness.

### Student activity

- ☐ Students attempt or retrieve something before receiving AI assistance.
- ☐ Students must evaluate, reject, revise, or justify parts of the AI response.
- ☐ Students identify what they checked, where they checked it, and what changed.

### Evidence of learning

- ☐ The task includes a period of work without AI access.
- ☐ Students retrieve, explain, apply, or transfer their learning afterward.
- ☐ I assess evidence of thinking as well as the polish of the final product.

**QUICK TEST**

**If the AI disappeared tomorrow, what would students still know or be able to do?**

If the answer is unclear, the task may be measuring access to assistance more than learning.

# READY-TO-USE CLASSROOM ACTIVITY

**Answer, Hint, Explain**

*Purpose: Help students compare answer-giving AI with AI used as a scaffold.  Time: 30–40 minutes.*

**01**

**Independent attempt**

Give students a moderately challenging problem, text, case, or question. Ask them to work independently for five minutes and record their initial answer, reasoning, and point of uncertainty.

**02**

**Request a complete answer**

Students ask AI for a complete answer. They mark one useful element, one claim or step that needs checking, and one part they still do not fully understand.

**03**

**Request a hint**

Give students a related task. They prompt: “Do not solve this for me. Give me one hint or ask me one question that will help me take the next step.” Allow up to two more hints.

**04**

**Explain the difference**

Students compare speed, effort, and reasoning in the two conditions. Ask which interaction is more likely to help them complete a similar task independently and why.

**05**

**Independent transfer**

Students complete a third related task without AI and identify which parts of the earlier interaction helped them.

**What to look for**

- Can students distinguish task completion from learning?
- Can they request assistance without surrendering the task?
- Can they evaluate a response and transfer the learning?

# PROFESSIONAL-LEARNING ACTIVITY

**Where did the thinking go?**

*Purpose: Examine whether an AI-supported task preserves the cognitive work connected to its learning goal. Time: 45–60 minutes.*

**01**

**Select a task**

Choose an existing assignment or activity in which students might use generative AI. Write down the intended learning outcome.

**02**

**Map the thinking**

List the cognitive actions students would normally perform: retrieve, interpret, compare, calculate, generate, organize, evaluate, revise, justify, or reflect.

**03**

**Add AI**

Identify which actions AI could perform. Label each one Protect, Support, or Offload. Protect the actions central to the goal; support selected actions; offload peripheral work.

**04**

**Redesign the task**

Add at least two safeguards: an independent attempt, hint-only prompting, a record of accepted and rejected suggestions, source verification, oral explanation, transfer problem, or delayed retrieval.

**05**

**Test the design**

Ask another group: “If a student completes this successfully, what evidence will we have that the student learned?”

# PRESENTATION-READY SUMMARY

**One slide for a staff meeting**

**WHEN DOES AI SUPPORT LEARNING?**

**Five conditions worth protecting**

AI follows an initial attempt; offers hints, questions, feedback, or contrasting examples; requires evaluation and revision; keeps the learner responsible for decisions; and is followed by independent retrieval or transfer.

**WHEN DOES AI REPLACE THINKING?**

**Five warning signs**

AI supplies complete answers too early; performs the reasoning targeted by the lesson; allows copying without evaluation; removes opportunities to persist through uncertainty; and produces a finished product that is never tested independently.

**THE DESIGN QUESTION**

**What is the student still doing cognitively?**

A polished product cannot answer this question. Look for visible evidence of retrieval, reasoning, evaluation, revision, and transfer.

## How Strong is The Evidence?

The studies cited in this brief provide different levels of evidence. Bastani, Fan, Barcaui, and Liu used randomized designs. Hawkins offers a small exploratory observation study. Lee and Gerlich report associations and self-reported experiences. Kosmyna is a preliminary multimethod manuscript. Lodge and Loble synthesize the wider literature.

*The convergence is useful, but the studies involve different populations, tasks, tools, and time frames. Their findings should guide careful design, not become a universal claim about all AI use.*

## FINAL TAKEAWAY

I do not think the answer is to keep AI outside the classroom. Students are already using it, and the capacity to work with AI thoughtfully is becoming an important part of education.

The more urgent task is to design its use around learning.

Students still need to retrieve knowledge, attempt difficult tasks, notice confusion, evaluate information, and explain their reasoning. They also need newer forms of judgment: knowing when to consult AI, how much assistance to request, what to verify, and when to put the tool aside.

**FIVE QUESTIONS**

**What remains after the assistance is removed?**

1. Can the student explain the idea?
2. Can they solve a related problem?
3. Can they identify a weak argument?
4. Can they remember the central concepts later?
5. Can they continue when an answer is not immediately available?

Those questions bring us closer to the kind of AI integration worth pursuing.

## Continue Learning with Me

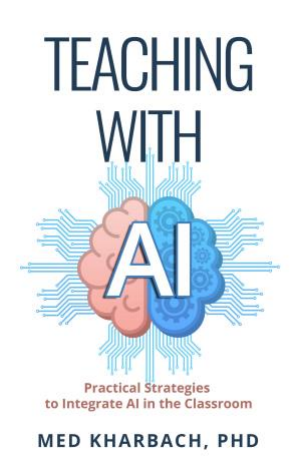


**GO DEEPER**

**Teaching with AI: Practical Strategies to Integrate AI in the Classroom**

A practical guide for educators who want to use AI purposefully while protecting student learning, judgment, and agency.

**Buy the book on Amazon →**

**RESEARCH DIGEST**

**Subscribe to the AI Research Digest**

**Research briefs and concise study reviews delivered to your inbox**

Join the AI and Education Research Digest for evidence-based insights on teaching, learning, assessment, and academic practice.

**Join the Research Digest →**

### AI Disclosure

I used AI tools while preparing this brief, including Claude Cowork and ChatGPT. Their role was limited to editing for language and style. I selected the sources, read them, and verified every claim and reference against the original documents. The analysis and conclusions are mine, as are any errors.

**Med Kharbach, PhD**

medkharbach.com/blog

## References

Barcaui, A. (2025). ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention. Social Sciences & Humanities Open. https://doi.org/10.1016/j.ssaho.2025.102287

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, 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

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. British Journal of Educational Technology, 56, 489–530. https://doi.org/10.1111/bjet.13544

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

Hawkins, B., Taylor-Griffiths, D., & Lodge, J. M. (2025). Summarise, elaborate, try again: Exploring the effect of feedback literacy on AI-enhanced essay writing. Assessment & Evaluation in Higher Education. https://doi.org/10.1080/02602938.2025.2492070

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). *Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task*. arXiv. https://doi.org/10.48550/arXiv.2506.08872

Lee, H.-P. H., 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*. https://doi.org/10.1145/3706598.3713778

Liu, G., Christian, B., Dumbalska, T., Bakker, M. A., & Dubey, R. (2026). AI assistance reduces persistence and hurts independent performance. *arXiv preprint arXiv:2604.04721*.

Lodge, J. M., & Loble, L. (2026). Artificial intelligence, cognitive offloading and implications for education. University of Technology Sydney. https://doi.org/10.71741/4pyxmbnjaq.31302475