

Who is *really* thinking here?



Generative AI does not settle next to education but inside it. As a conversation partner, co-writer and suggestion machine. Rules say what is allowed and training teaches how the tools work, but between the two lies a space we rarely have words for. That is exactly where learning happens. The HEAR framework gives that middle ground four names and one conversation.



Pascal Mariany

Senior lecturer at HU University of Applied Sciences Utrecht, founder of CivAI, the European sovereign AI platform behind EduGPT, and founder of Onderwijs Orakel. This article is based on chapter 6 of his book *De kunst van het samenleven met AI* (Van Duuren Media, 2026). The English edition, *The Art of Living with AI*, is due in October 2026.

Where the friction is

The friction is not in the technology. It is in the moment a student receives an answer before working out the question, and no longer knows whether he has been helped or overtaken. When feedback reads smoothly but the conversation about choices disappears. When suggestions sound logical and slowly steer thinking and action, without anyone consciously choosing that.

These are not technical shortcomings. They are shifts in meaning. And meaning does not arise in a model, but in the way people work with that model. In language, in timing, in intention. In the moment someone decides to take over here, and to be led there.

As a lecturer I see that moment every day. As a founder I speak with school leaders looking for a grip on it. At both tables the same thing is often missing, namely words for what happens in the interaction itself.


Why rules and buttons are not enough

Institutions understandably look for something to hold on to. They draw up rules, make lists of approved tools, invest in detection or decide not to lay anything down for now. All of these responses are understandable, and yet none of them touches the core problem.

Rules say something about what is allowed, but little about what happens. Skills training teaches people how to

use AI, but not when they would do better to leave it alone. Detection shifts the conversation to control, while the problem sits in meaning and responsibility. And leaving everything open leads to invisible use and strategic behavior.

Research from the past two years shows how real the shift is. Gerlich (2025) found among 666 participants a clear link between intensive AI use and lower scores on critical thinking, with cognitive offloading as the explanatory link. Lee and colleagues (2025) saw among 319 knowledge workers that more confidence in AI went together with less critical thinking of their own, while confidence in their own abilities did the opposite.



The most telling picture comes from a field experiment by Bastani and colleagues (2025) among nearly a thousand high school students. With a standard chatbot, students performed better during practice, but worse on the test without AI than the control group. A variant that only gave hints instead of answers largely removed that disadvantage. The design of the interaction, not the tool, determined the learning effect.

Zhai, Wibowo and Li (2024) sum up the literature in one warning. Whoever adopts AI answers unquestioningly loses not only knowledge, but also the motivation to think.

The risk is not that students know less, but that they think less.

FROM THE ART OF LIVING WITH AI

Four layers, one conversation

HEAR grew out of that observation. Not as an answer to how we should regulate AI, but to how we can understand what happens when people work with it. It is not a method you implement or tick off, but a way of looking. The letters stand for Human Expressive AI Reasoning. Two layers sit with the human being, two in the interaction with the system (figure 1). The human being stands at the beginning and at the end. In between, much happens that we usually do not see.

H LAYER 1 Human Intent

Before a prompt is typed, there is always something else. An intention. A learning goal, a decision to be made, a search for grip on complexity. The same question to AI can spring from very different intentions. Someone wants to test whether he understands the core, save time, mask uncertainty or outsource responsibility. AI cannot tell the difference. The human being has to.

Research on self-regulated learning shows why that matters. Learning becomes more effective when learners make their goals explicit and reflect on their own role (Zimmerman, 2002; Panadero, 2017). Generative AI increases the risk that this step is skipped, because it answers quickly and smooths over uncertainty. That is why Human Intent is a cognitive anchor. A moment of slowing down before speeding up.

What do I actually want to learn, achieve or decide here, and which responsibility stays with me?

THE FOUR LAYERS OF HUMAN INTENT

- Goal intention**
What do I want to achieve?
Understand, practice, design, decide or articulate.
- Learning intention**
What do I need to develop myself?
Understanding, skill, judgment, reflection.
- Role intention**
Am I acting here as a learner, a designer, an assessor or a coach?
- Responsibility intention**
What remains my responsibility, even with AI? Accuracy, quality, ethics, the decision.

E LAYER 2 Expressive Interaction

Language is never a neutral conduit. Generative AI is an expressive system that presents answers in fluent sentences and a tone that suggests certainty. Users judge fluent output as correct more quickly, even when substantiation is missing (Kasneci et al., 2023). Bender and colleagues (2021) aptly called such a system a parrot that produces form without meaning.

For learning, that is a problem. Fluent language creates the impression of understanding, and a confident explanation makes uncertainty disappear, even when that uncertainty was pedagogically valuable. Expressive Interaction invites us to take language seriously as a didactic factor.

Do I follow because the content is right, or because the form convinces?

Not everything that is clear has been learned. Not everything that sounds logical has been understood.

FROM THE ART OF LIVING WITH AI

A LAYER 3 AI Mediation

Between question and answer there is no direct line. There is a system that predicts what is probably fitting, not what is true, complete or pedagogically desirable. What appears is always a filtered possibility, one proposal from a much larger field. What is not mentioned disappears from view. What is smoothly formulated gains weight.

Silvola and colleagues (2025) analyzed how students wrote with generative AI and showed that the system co-steers their sensemaking. Students adjusted text, questions and approach to what AI offered. Lines of reasoning appear as finished products, considerations have already been made implicitly, counter-arguments remain invisible. Learning thus shifts from active construction to passive interpretation.

AI Mediation is the hinge of HEAR. As long as the mediation remains invisible, AI seems self-evident. As soon as it becomes visible, room emerges to slow down, open up alternatives and choose anew.

Which choices had already been made before I started thinking?

R LAYER 4 Reasoned Responsibility

Who carries the responsibility when AI thinks along, co-writes and sometimes co-decides? AI is not a moral subject. It

can present plausible steps, but it cannot carry what human action always carries, namely responsibility for consequences. Weizenbaum (1976) already warned of the quiet shift from judgment to calculation. UNESCO (2023) states that human oversight must not disappear through generative AI, and the European AI Act has made AI literacy an obligation for anyone deploying AI since February 2025.

In practice the problem rarely arises from one big mistake. First AI takes over a formulation, then a line of argument, then a structure. Reasoned Responsibility restores ownership of judgment, grounded in traceable considerations and not in control or suspicion. It makes the difference between AI as a companion and AI as a substitute.

Can I explain why something is right, which assumptions I made and where my own judgment lies?

From task to learning

What changes when you look through HEAR? The central question shifts from 'is AI allowed?' to 'what thinking do I want to see?' (figure 2).

For students this means a shift from task-oriented to learning-oriented use. Whoever asks AI to write an essay asks for an end product. Whoever asks AI to question their own argumentation uses AI as a mirror for thinking. The difference is not in the quality of the output, but in who thinks, who chooses and who learns.

For teachers the attention shifts from detection to guided judgment. Not whether AI was used, but whether the

thinking is visible and the responsibility properly placed. That calls for assignments you design rather than secure, and for moments when students have to put their reasoning into words.

For teams and curriculum designers HEAR is a design parameter. Not every learning goal suits the same degree of AI support. In formative feedback AI can be valuable because the process is central. In summative assessment the same use can undermine the intention. Crolla and colleagues (2026) show why that conversation is urgent. Strong students use AI to extend their thinking, vulnerable students outsource it. One rule for everyone produces unequal outcomes.

For school leaders and policymakers Human Intent is a guideline that goes beyond tool choice. Policy that starts from intention is less prone to hype and more resistant to technological change. Swiss research by Consoli and Petko (2025) among more than 2,300 students is encouraging. A critical-reflective approach in the classroom predicted both students' confidence and their sense of responsibility around AI. Responsibility does not grow from use by itself, but from didactics.

Intention without language awareness remains vulnerable. Language without intention becomes steering.

FROM THE ART OF LIVING WITH AI

ROLE	FROM TASK-ORIENTED	TO LEARNING-ORIENTED
Student	Asks AI for the end product and adopts the wording.	Has AI question their own argumentation and rewrites the proposal in their own words.
Teacher	Asks whether AI was used and looks for proof.	Designs assignments that make thinking visible and asks about intention and considerations.
Team	Chooses a tool list that applies equally to every subject.	Decides per learning goal which role AI may play and where AI stays deliberately absent.
Leadership	Decides which AI tool is permitted.	Starts from intention and names which professional judgment it wants to strengthen.

Figure 2. The same four layers, four roles. HEAR shifts the central question from 'is AI allowed?' to 'what thinking do I want to see?'. After Mariany (2026).



EduGPT

HEAR in practice

Especially in a time of AI, I believe teachers are designers, not operators of a tool. With that conviction we build AI companions within EduGPT that are designed on HEAR. The Feedback Companion does not take over the teacher's judgment, but makes choices explicit and keeps the teacher's intention in the lead.

The Curriculum Companion does not assess a curriculum, but puts design questions on the table. Where is independence assumed without being prepared? Where does AI add something and where should AI stay deliberately absent? Human Intent then sits with the team, and Reasoned Responsibility becomes collective.

Wang and colleagues (2026) describe a similar move with their IDEA framework, which builds in metacognitive regulation against what they call 'performance without learning'. The thinking stays where it belongs, with people.

Six starting points for teachers and teams

1 Start with the intention question

Have students write down in one sentence what they want to learn or decide before the first prompt.

2 Do the mirror test

Have an AI answer rewritten in the student's own words. What remains is understanding, what falls away was form.

3 Design for slowing down

Build in moments when AI is not yet involved, and explain why.

4 Make language a learning goal

Discuss where a text convinces by content and where by tone. Fluent is not the same as true.

5 Assess the account, not the use

Ask what the intention was, what AI did and what the student weighed up, not whether AI was used.

6 Make it a team agreement


Decide per learning goal which role AI may play and where it stays deliberately absent.

Never done learning

HEAR gives no answers. It gives a shared point of reference for having the same conversation, in the classroom, in the team and at the board table. Not about perfect answers, but about people who learn to think, to choose and to carry what they do.

Only humans can have intentions. HEAR begins and ends there. I am never done learning and would love to continue this conversation. Where in your teaching is the thinking still visible, and where has it quietly shifted to the system?

Want to respond or talk further?
pascal@civai.eu



Read more

Mariany, P. (2026). *The Art of Living with AI. On learning, teaching and humanity in the age of AI*. Van Duuren Media.

**ENGLISH EDITION DUE
OCTOBER 2026**

Sources

- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakci, 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. doi.org/10.1073/pnas.2422633122
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). doi.org/10.1145/3442188.3445922
- Consoli, T., & Petko, D. (2025). Which educational approaches predict students' generative AI confidence and responsibility? *Computers and Education: Artificial Intelligence*, 9, 100431. doi.org/10.1016/j.caeai.2025.100431
- Crolla, K., Xia, X., & Jiang, Y. (2026). AI-mediated cognitive divergence in built-environment education: Evidence from a mixed-methods study. *Computers and Education: Artificial Intelligence*, 11, 100665. doi.org/10.1016/j.caeai.2026.100665
- European Union. (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 4. *Official Journal of the European Union*. eur-lex.europa.eu/eli/reg/2024/1689/oj
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 6. doi.org/10.3390/soc15010006
- Kasnezi, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., ... Kasnezi, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. doi.org/10.1016/j.lindif.2023.102274
- Lee, H.-P., 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. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. doi.org/10.1145/3706598.3713778
- Mariany, P. (2026). *De kunst van het samenleven met AI. Over leren, onderwijzen en menselijkheid in het tijdperk van AI*. Van Duuren Media. English edition *The Art of Living with AI* forthcoming, October 2026.
- Panadero, E. (2017). A review of self-regulated learning: Six models and four directions for research. *Frontiers in Psychology*, 8, 422. doi.org/10.3389/fpsyg.2017.00422
- Silvola, A., Kajamaa, A., Merikko, J., & Muukkonen, H. (2025). AI-mediated sensemaking in higher education students' learning processes: Tensions, sensemaking practices, and AI-assigned purposes. *British Journal of Educational Technology*. doi.org/10.1111/bjet.13606
- UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO. unesco.org/en/articles/guidance-generative-ai-education-and-research
- Wang, X., Zheng, Z., Zhang, J., Hou, X., & Zhu, Z. (2026). The IDEA framework for metacognitively regulated GenAI use in higher education: Development and exploratory pilot evidence. *Computers and Education: Artificial Intelligence*, 11, 100657. doi.org/10.1016/j.caeai.2026.100657
- Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. W. H. Freeman.
- Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments*, 11, 28. doi.org/10.1186/s40561-024-00316-7
- Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, 41(2), 64–70. doi.org/10.1207/s15430421tip4102_2

Images for illustration, generated with AI (OpenAI Image 2.5 via Azure AI Foundry EU). Figures 1 and 2 are the author's own. © 2026 Pascal Mariany, Civai B.V. Share with attribution.