Adv Mobile Learn Educ Res, 2026, 6(2): 1940-1951
DOI: 10.25082/AMLER.2026.02.007

PERSPECTIVE

A Teacher's Concerns about AI in Primary School: Critical Thinking, Personal Data, and the Fear of Cognitive Dependence

Foteini Kareli

10th Primary School of Aigio, Aigio 25100, Greece

Check for updates

Correspondence to: Foteini Kareli, 10th Primary School of Aigio, Aigio 25100, Greece;
Email: kareli.fot@gmail.com

Received: August 6, 2026;
Accepted: November 7, 2026;
Published: November 11, 2026.

Citation: Kareli, F. (2026). A Teacher's Concerns about AI in Primary School: Critical Thinking, Personal Data and the Fear of Cognitive Dependence. *Advances in Mobile Learning Educational Research*, 6(2), 1940-1951.
https://doi.org/10.25082/AMLER.2026.02.007

Copyright: © 2026 Foteini Kareli. This is an open access article distributed under the terms of the Creative Commons Attribution-Noncommercial 4.0 International License, which permits all noncommercial use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: Generative artificial intelligence (AI) is already part of pupils' everyday lives and of school life. Written from the perspective of a primary school teacher and principal, this perspective paper examines three interrelated concerns of the author herself. These concerns arose from a wish to protect her pupils and led her to study the topic for her own professional development and to help her colleagues make teaching safer and better. They concern developing critical thinking towards AI-generated answers, protecting the personal data of minors, and the fear that excessive assistance from AI will weaken pupils' thinking. Drawing on a narrative, non-systematic review of recent research on cognitive offloading, the European and Greek regulatory framework, and illustrative examples from the author's own teaching practice and lesson scenarios for primary school, it argues that the fear of cognitive dependence is a plausible concern rather than an established finding for primary school pupils. The available evidence, which comes mainly from adolescents and adults, suggests that the form of AI support (whether it replaces the learner's effort or scaffolds it) may be an important factor. AI is regarded throughout as a valuable tool; the risk lies mainly in the way it is used. The classroom examples come from ordinary lessons whose only purpose was to teach pupils to question AI answers and to protect themselves when using AI. They were not planned as research, no research data were collected, and they are offered as the author's practitioner reflections, not as findings. The author's classroom experience illustrates that such activities can be implemented in ordinary primary school lessons, but their effects on pupils' critical thinking have not yet been evaluated. The paper relates the three concerns to long-standing educational principles, in particular Socratic questioning, Dewey's reflective thought and Vygotskian scaffolding. On this basis, it proposes a simple, as yet untested three-step classroom protocol ("think – ask – check"), age-appropriate data protection rules, and a role for school leadership, and outlines questions for future research.

Keywords: cognitive offloading, AI literacy, children's privacy, GDPR, scaffolding, primary education, school leadership

1 Introduction

A ten-year-old child today can ask a language model to solve a mathematics problem, write an essay, or explain why it rains. The answer will arrive within seconds; it will be well written and, most of the time, correct. It is precisely this ease that raises the questions currently troubling those who teach in primary school. If the answer is wrong, will the child notice? What information about themselves and their family do children leave behind when they converse with such a system? And, most worrying of all: if the tool thinks on the child's behalf, will the child ever learn to think independently?

For anyone teaching in primary school today, these questions are not theoretical. This paper is written from the perspective of a teacher and primary school principal who first encountered them in the classroom, in children's questions and in parents' concerns. It arose from the author's fear of what might happen to her pupils and from the need to protect them. The paper did not begin as a research project. The classroom activities described in it were ordinary lessons, planned with a single aim: to help the author's pupils see that AI can be wrong and to teach them to protect themselves when they use it. Only later, having learned a great deal from these lessons and from the reading that accompanied them, did the author decide to summarise what she had learned and share it with fellow teachers and with anyone else interested. The study of the literature was initially undertaken for her own professional development, out of

love for her work. The aim of the paper is to share what she has learned with fellow teachers, so that together we can make teaching safer and better.

These questions are not new in the history of education. Similar concerns accompanied the pocket calculator, the search engine and the mobile phone. Generative AI, however, differs in one crucial respect: it does not merely take over a mechanical function, such as a calculation, but can perform the entire work of thinking, from synthesising information to formulating conclusions (Kasneci et al., 2023). Moreover, it does so with a linguistic fluency that lends its answers authority, even when they are wrong (Ji et al., 2023).

From the outset, the author views AI as a valuable tool and capable assistant, both for the teacher and for the pupil. The concerns that follow are not about the tool itself but mainly about the way we use it. When AI is used inappropriately in teaching, the responsibility lies primarily with the people who decide how it is used. The analogy of the car is illuminating. A very good car protects its passengers even at very high speeds thanks to its construction. But if the driver loses control, they will lose it at high speed, and the consequences can be serious. The more powerful a tool is, the greater the responsibility of the person handling it.

Teachers' own perspectives are an important part of this picture. In a qualitative study with early childhood teachers who had been trained in the use of ChatGPT, Uğraş et al. (2024) found that teachers valued the tool in STEM activities for instant feedback, personalised content and the encouragement of creativity, but also reported technical difficulties and concerns about security, and emphasised that pupils need information-literacy training and teachers need pedagogical support. In preschool education, Katsampoxaki-Hodgetts et al. (2025) analysed sixty AI-enhanced teaching scenarios developed by teachers during a national training programme and suggested that a systematic, multimodal and culturally sensitive approach can support the development of early critical literacy. The present paper adds the perspective of a primary school teacher and principal, with a particular focus on critical thinking, personal data and cognitive dependence.

The European and Greek regulatory frameworks (see Section 3.1) define what is permitted and what must be protected. On their own, however, they do not answer the pedagogical question: how do we help a child of eight or ten to think *with* the machine without handing their thinking over to it? This question belongs to us, the teachers, and it is to this question that the present paper seeks to contribute.

The paper is a perspective piece. It is based on a narrative, non-systematic review of the literature and on the author's reflection on her own teaching practice; it does not report a research study. Each concern is examined first in the light of the literature and then through illustrative examples from the author's own teaching practice: lessons she has taught in a primary school and lesson scenarios designed for the classroom. Related classroom material, on AI as an assistive technology for pupils with learning difficulties and ADHD, has been accepted for oral presentation at an educational conference and has not yet been published (Kareli, 2026); In each case it is made clear whether an activity was implemented or is proposed. The classroom examples are practitioner accounts. The lessons were not planned as research, no data were collected for research purposes, and the examples are used to illustrate the arguments, not to provide evidence for them. The three questions are: (a) how critical thinking towards AI can be cultivated in primary school children; (b) what obligations and risks the use of AI entails for pupils' personal data; and (c) what research says about the fear that AI weakens thinking. The position argued is that the three issues are connected: a child who has learned to question the machine is also a child who knows what not to entrust to it, and when not to delegate their own thinking to it.

1.1 Approach to the literature

The review is narrative and was not conducted according to a systematic protocol: no predefined database search, search terms, date limits, formal inclusion criteria or study-selection procedure were used. Sources were selected purposively by the author because they directly address one of the three concerns of the paper or are widely cited in the field. They comprise foundational works on critical thinking, scaffolding and cognitive offloading; the European and Greek regulatory texts in force at the time of writing; international policy guidance (UNESCO, UNICEF); and recent empirical studies on generative AI and thinking. Because very few studies concern primary school pupils, most of the empirical evidence discussed comes from adolescents and adults, and its transfer to younger children is treated throughout as an open question. The selection is therefore illustrative rather than exhaustive, and the conclusions drawn from it are tentative.

2 Concern 1: Will children recognise machine fallibility?

2.1 From critical thinking to AI literacy

Critical thinking has been defined as purposeful, self-regulatory judgement that results in interpretation, analysis, evaluation and inference, together with an explanation of the evidential and criteriological considerations on which that judgement is based (Facione, 1990). In the context of AI, this skill falls within the broader concept of AI literacy, which Long and Magerko (2020) describe as a set of competencies that enable people to critically evaluate AI technologies, communicate and collaborate effectively with AI, and use AI as a tool online, at home and in the workplace. The UNESCO AI competency framework for students (Miao & Shiohira, 2024) places a human-centred mindset, the ethics of AI and an understanding of AI techniques at its core, while Regulation (EU) 2024/1689 on artificial intelligence already requires deployers of AI systems to take measures to ensure a sufficient level of AI literacy among the people who operate them (Regulation, 2024).

For a primary school child, critical thinking towards AI can be described more simply as three habits: knowing that the machine can be wrong, knowing how to check an answer, and being able to explain why one accepts or rejects it. None of the three is self-evident. Language models occasionally produce answers that are persuasively worded but inaccurate, a phenomenon referred to in the literature as "hallucination" (Ji et al., 2023). Furthermore, research on human interaction with automated systems has shown that people tend to over-trust a system's suggestions and reduce their own monitoring, a phenomenon known as automation bias (Parasuraman & Manzey, 2010). Primary school pupils may also be vulnerable to automation bias, although direct evidence concerning children in this age group remains limited.

2.2 Illustrations from classroom practice

A critical attitude towards AI is hard to teach through sermons. Guided discovery, in which the child sees the machine err with their own eyes, offers a more concrete alternative, although its effectiveness for this purpose has not been tested. In a digital-safety lesson at the school where the author works, the teacher asked a widely used generative AI tool, in front of the class and on the interactive whiteboard, to present the content of a specific page of the printed school textbook that the pupils had in their hands. The AI presented a different page from the one requested, because it had no access to the pupils' printed book, whereas the official digital edition of the same book displays the page correctly. What was pedagogically significant was that, despite this lack of access, the AI did not reply that it did not know. It gave a wrong answer, and did so with confidence. The pupils were asked to compare the projection with their printed book. The teacher did not immediately confirm their observations but let the children defend their position against the machine's "certainty".

The activity was carried out as part of ordinary lessons in three classes of different grades (Grades 3, 4 and 6). Its only purpose was to teach the pupils that an AI answer must be checked before it is trusted, so that they could protect themselves when using such tools. It was not planned as a research study. No notes were kept during the lesson and no material produced by pupils was collected, because no later use was intended. The account that follows is based on notes the author wrote from memory afterwards, when she decided to summarise her experience, and on her later reflection, as the class teacher, on what she observed. No numerical results are reported; it is not a measurement, and it is not intended to support comparisons between grades.

Three observations from this practice are relevant to the paper, with the caveat that they are the impressions of a single teacher. First, even in the Grade 3 class, some pupils noticed the discrepancy before the teacher prompted them. This suggests that questioning an AI answer is not necessarily beyond the reach of younger pupils when the task is well supported, although it says nothing about how widespread this ability is. Second, the printed book served as a fixed point of reference: the existence of an independent, reliable source made comparison, and therefore questioning, possible. The lesson for the children is not that paper is more reliable than the screen, since the official digital edition of the book showed the same correct page. It is that an answer is checked against a reliable source, whether printed or digital. Third, the justifications the pupils offered differed. Some simply noted the difference from the book, whereas others also offered an explanation for the cause of the error, namely that the AI did not have access to the pupils' book. Because the classes differed not only in age but also in prior experience, composition and the amount of support the teacher provided, these differences

cannot be attributed to age or development. They are mentioned only because they illustrate the role that teacher support can play, in line with the concept of the zone of proximal development (Vygotsky, 1978) and with the idea of scaffolding, that is, temporary support that the adult gradually withdraws as the child's autonomy grows (Wood et al., 1976).

In other lessons with Grade 6 pupils (age 11–12), critical thinking was directed not only at the accuracy but also at the nature of AI answers. In an activity aimed at understanding how AI works, a language model was asked to describe how a child would feel if, after hours of playing at the beach during a heatwave, their mother treated them to an ice cream. The answer was strikingly vivid: it described the burning sun, the salt on the lips, the coolness, the joy and relief, even the choice of flavour. The pupils then asked, in effect, how the system could know all this and whether it had ever eaten an ice cream. The system itself replied that it has no senses, emotions or experiences and that it composes its answers from patterns and associations in large volumes of text. This moment functioned as a cognitive conflict: the pupils had before them a text that "looked" experiential, but they now knew that it rested on no experience at all. The contradiction led them to distinguish between a text that seems true and a text grounded in experience or evidence, a distinction which, in our view, lies at the heart of critical thinking towards generative AI.

In the same series of lessons, the pupils were first asked to describe a doctor themselves. Their descriptions were personal: they talked about their own paediatrician, the way he spoke to them, their fear before an examination. When the same question was given to an AI system, the answer was a generic, stereotypical description without any personal element. The comparison opened up three questions in the class: whether the machine had answered what we had actually asked, whether the question had been worded correctly, and how we can check an answer. The discussion culminated when the children wondered whether they would trust a diagnosis or treatment proposed by AI. In this way, critical thinking was transferred from a school example to a question with real consequences for their lives.

The most complete application was a four-hour lesson scenario for Grade 6 on a hypothetical water shortage in the town. The pupils, acting as a "team of advisers to the municipality", searched for information on the internet and with AI tools, checked whether the proposals they obtained were applicable to their own area, put together their own solution and defended it in front of the class, which acted as the "municipal council". The crucial element was assessment. What was graded was not how good the final proposal was, but whether the pupils could explain why they had rejected some of the AI's ideas and why they had kept others. In this way, AI became an occasion for judgement rather than a source of ready-made answers.

These examples come from the author's own lessons in a single school and are offered as illustrations, not as findings. They were not evaluated systematically and cannot be generalised. They do suggest, however, that activities of this kind can be integrated into ordinary lessons with simple means; whether they cultivate a lasting critical attitude in pupils is a question for future research.

3 Concern 2: What do children leave behind?

3.1 The regulatory framework

The General Data Protection Regulation (GDPR, Regulation (EU) 2016/679) recognises that children merit specific protection, as they may be less aware of the risks involved and of their rights (Recital 38) (Regulation, 2016). It sets data minimisation as a core principle, i.e., processing only what is necessary for the purpose (Article 5(1)(c)), and requires data protection by design (Article 25). In Greece, when an information society service, such as an AI application, is offered directly to a minor and the processing of their data is based on consent, the child may consent alone only from the age of 15. For younger children, the consent of their legal representative is required (Law, 2019; Regulation, 2016). In other words, no primary school pupil can validly consent on their own to the processing of their data by such a service.

Regulation (EU) 2024/1689 on artificial intelligence (Regulation, 2024) adds two provisions of direct relevance to schools. It prohibits the use of AI systems to infer the emotions of natural persons in education institutions, except for medical or safety reasons (Article 5(1)(f)). It also classifies as high-risk, among others, systems used to evaluate learning outcomes or to monitor pupils during tests (Annex III, point 3). At national level, Joint Ministerial Decision 55909/Δ6/2026 Ministerial Decision (2026), which applies to secondary schools, provides that AI is used on a voluntary basis and only after training (Article 2). It prohibits the entry of identifying data into AI systems, such as names, photographs, videos, voice samples, contact

details, grades, disciplinary data, health data and information on special educational needs, and permits only neutral, non-identifying, synthetic or anonymised data (Article 16). It also prohibits fully automated assessment decisions and the creation of individual profiles of pupils or teachers (Article 4), and requires human oversight at every stage (Article 7). Although the decision does not apply to primary schools, its principles constitute, in our view, the minimum reasonable standard for younger children as well.

3.2 Risks in practice

Risks to pupils' data do not come only from the children themselves. A teacher who asks an AI tool to adapt exercises "for George P. in Grade 4, who has an ADHD diagnosis" (Attention-Deficit/Hyperactivity Disorder) has already transferred to a third-party provider a full name, a class and health data, that is, a special category of data under the GDPR. The same result, namely exercises adapted for "a Grade 4 pupil with attention difficulties", can be achieved without any identifying information. The principle of data minimisation thus becomes a very concrete professional habit.

A second risk is linked to the capabilities of generative AI itself. From a few seconds of audio or a few photographs, such as those often posted on school websites or social media, it is now technically feasible to clone a voice or create fake audiovisual material featuring a child's face. UNICEF (2021) points out that children are particularly exposed to such risks and that protecting their privacy should be a design principle of any AI system that concerns them. Livingstone, Livingstone et al. (2019) also show that children understand the data they share knowingly better than data that is collected silently or inferred about them by platforms. This reinforces the responsibility of adults: parents, teachers and school leaders.

3.3 Data protection as a lesson

Data protection can also become a subject of teaching. The author has designed a lesson scenario for Grade 6 (age 11–12) with digital-safety activities. The same activities also appear, with a different focus, in a conference paper on AI as an assistive technology for the inclusion of pupils with learning difficulties and ADHD (Kareli, 2026), which has not yet been published. They are summarised here so that the present paper can be read on its own, and they are considered only in relation to its questions. Their first logic is data minimisation. When children compare two versions of the same request to an AI tool, one containing the personal details of a fictitious pupil and one without them, the intended point is that they can see that the answer is no worse when the details are missing. This conclusion can become a class rule, written on a privacy card on the classroom wall, which sets out what information never goes into a digital tool. At the same time, the teacher acts as a model: the children see the teacher phrasing their own requests without exposing anyone.

Their second logic is trust. Two further activities of the same scenario concern online scams and the digital footprint. In the first, the teacher uses AI to create realistic but harmless versions of a fake "you have won a prize" pop-up message, and the class discusses how to recognise them. In the second, short AI-generated dialogues prompt pupils to pause and think before sending or posting something. These activities (Kareli, 2026) aim not only to help children recognise the risk, but also to make sure they know that, if they make a mistake, they can tell an adult without fear of punishment. From the point of view of child protection, this trust is worth as much as knowledge. Activities of this kind often grow out of everyday school life rather than out of a plan. An incident familiar to many teachers is a pupil posting an offensive comment about a classmate on a social media platform. Deleting the comment does not undo the hurt it has already caused, but such an incident, discussed later in general terms and without identifying anyone, can become a lesson that helps pupils avoid repeating the same mistake, both as children and in later life. Generative AI adds a new dimension here. A comment, a photograph or a recording that a child leaves online may persist even after it has been deleted, and it can now also be edited and reused in ways the child cannot imagine. Before posting anything, therefore, children should learn to ask not only what others will think today, but also whether they would want to come across it again years from now.

4 Concern 3: Will children develop independent thought?

4.1 Cognitive offloading: an old phenomenon on a new scale

The fear that AI will make children "mentally lazy" is linked to a well-studied phenomenon, cognitive offloading, i.e., the use of external means to reduce the cognitive demands of a task

(Risko & Gilbert, 2016). Offloading is not negative in itself. The shopping list, the diary and the calculator free up resources for more demanding tasks. The problem arises when the external aid takes over precisely the process we want to learn.

Research predating generative AI had already shown that expected access to information affects memory. When participants believed that information would remain available on a computer, they remembered less of the information itself and more about where it was stored (Sparrow et al., 2011). Similarly, habitual use of GPS navigation has been associated with poorer spatial memory when people must find their way on their own(Dahmani & Bohbot, 2020). These studies, conducted with adults and in domains other than school learning, suggest that relying on an external aid can be associated with weaker performance when the aid is removed. They concern memory and navigation, not the learning of school skills by children, and are cited here only as background to the hypothesis examined below.

4.2 What recent research on generative AI says

In reading this research, four outcomes must be kept distinct: immediate task performance (how well a learner completes a task, with or without AI); learning retention (what remains when the aid is removed); critical-thinking disposition (the tendency to question and evaluate information and arguments); and cognitive offloading (the delegation of mental processing to an external aid). The studies discussed below measure different ones of these outcomes, often in different ways, and a finding about one cannot be read as evidence about another. Research on generative AI is new, and for primary school children it is almost non-existent. The first data from adolescents and adults are, however, instructive. In a field experiment with nearly 1,000 high school students (grades 9–11, approximately aged 14–17) in Turkey, unrestricted access to a general-purpose language model increased performance on practice exercises by 48% compared with the control group. When the assistance was removed, however, the same students performed 17% worse on the exam than those who had not used AI. By contrast, a version of the tool designed as a "tutor", giving hints instead of solutions, improved practice performance even more without harming exam performance (Bastani et al., 2025). This study concerns task performance and retention in high school mathematics; it did not measure critical thinking. Its finding is nevertheless useful for teachers: in this setting, the harm was associated with AI that gave ready-made answers, not with AI as such.

In a study with 666 adults in the United Kingdom, frequent use of AI tools was negatively correlated with scores on a critical thinking scale, with cognitive offloading mediating the relationship; dependence was stronger among younger participants (Gerlich, 2025). Being correlational and based largely on self-report, this study cannot show that AI use causes weaker critical thinking. In a survey of 319 knowledge workers, the greater the participants' confidence in AI, the less critical thinking they reported exercising, whereas those with greater confidence in their own abilities reported exercising more (Lee et al., 2025); here too, the outcome is self-reported effort, not measured ability. A systematic review concluded that over-reliance on AI dialogue systems may negatively affect students' decision-making, critical thinking and analytical reasoning (Zhai et al., 2024). Finally, an electroencephalography study of 54 adults writing essays recorded lower neural connectivity and weaker recall of their own text among those who used a language model (Kosmyna et al., 2025). This study, however, is small and, at the time of writing, had not been published in a peer-reviewed journal, so caution is needed in its interpretation; it concerns adults and cannot be generalised to children.

Overall, the data support the fear of cognitive weakening as a plausible hypothesis rather than an established fact, and they do not establish it for primary school pupils at all. Most studies are correlational or based on self-reports, concern adults or adolescents, and cover short periods of time. In the one experimental study that compared different forms of use (Bastani et al., 2025), the outcome depended on the manner of use: AI that supplied solutions was associated with lower exam performance, whereas AI that offered hints was not. Whether this pattern holds for younger children, and for critical-thinking disposition rather than performance, is not known.

4.3 Why primary school is particularly vulnerable

There are theoretical reasons to think that the concern may be greater in primary school, although direct evidence is lacking: an adult who delegates a text to AI has already learned to write, whereas a primary school child is precisely at the stage of building basic skills. Research on learning has shown that certain "desirable difficulties", such as the effort of retrieval or struggling with a problem before seeing the solution, enhance long-term learning, even if they worsen immediate performance (Bjork, 1994). A tool that removes every difficulty may make

the day's homework easier and learning poorer.

The goal, therefore, is not to keep AI out of school. Pupils already use it at home, and a ban would leave that use without guidance. Miao and Holmes (Miao & Holmes, 2023), in UNESCO's guidance on generative AI, propose a minimum age of 13 for the independent use of generative AI tools in the classroom. This means that in primary school, AI is used mainly by the teacher, in front of and together with the children, as an object of discussion and not as the pupil's personal assistant.

4.4 From the classroom: scaffold or crutch?

The image of the scaffold rests on Vygotsky's theory (Vygotsky, 1978). According to Vygotsky, with the help of an adult or a more capable peer, a child can achieve more than they can alone. The distance between the two is the zone of proximal development, and it is there that teaching has its greatest value. The term scaffolding for this kind of assistance was subsequently introduced by Wood, Bruner and Ross (Wood et al., 1976). The scaffold supports the child as long as necessary and is gradually removed, like the scaffolding of a building once it can stand on its own. A crutch, by contrast, is never removed. If AI assistance is not withdrawn, the child does not learn to walk without it.

The distinction between AI that replaces thinking and AI that supports it can be seen clearly in one of the most common requests teachers make to AI: turning a complex instruction into a card of sequential steps. This use is also discussed, for a Grade 6 class including pupils with learning difficulties and ADHD, in the author's unpublished conference paper mentioned in Section 3.3 (Kareli, 2026). For a child with limited working memory, an instruction that asks at once for reading, identifying arguments and writing a paragraph often leads to confusion and giving up, whereas one step at a time allows the child to get started. In this case, AI does not think on the child's behalf. It removes the unnecessary load of organising the task and leaves the essential work intact: finding the arguments and formulating an opinion of one's own.

The same technique, however, also shows how fine the line is. When AI breaks a mathematics problem into steps, for example the calculation of a discount, it often includes the operation itself and its result in the final step. A card of this kind also appears in the paper mentioned above (Kareli, 2026), where it is discussed as support for pupils with ADHD and dyscalculia. Seen through the lens of the present paper, however, this final step is exactly the point at which the scaffold becomes a crutch. If the card stops before the final calculation, the pupil carries out the operation, checks whether the result is reasonable and experiences the success as their own. The change is small, but it corresponds to the difference that Bastani et al. (2025) found between the tool that gave solutions and the one that gave only hints. The conclusion for teachers is practical: any material that AI produces for pupils needs to be checked with the question "which part of the thinking does this material leave to the child?"

5 Educational principles still applicable

Many of the answers to the three concerns were formulated long before AI. Three connections are particularly direct. Socrates, in Plato's *Meno*, does not give the young slave the solution to the geometry problem. He leads him through questions to find it himself, after first passing through the confusion of error (Plato, *Meno*, 82b–85b). A tool that gives the answer immediately tends to do the opposite. The finding of Bastani et al. (2025), that AI giving hints instead of solutions did not harm exam performance, is consistent with the logic of Socratic questioning.

Dewey (1910) defined reflective thought as the active, persistent and careful consideration of any belief in the light of the grounds that support it, and regarded suspended judgement as the essence of critical thinking. This is what the "check" step asks of the child: not to accept an answer merely because it is well written.

Vygotsky (1978) showed that adult help is most valuable when it operates within the child's zone of proximal development and is gradually withdrawn. This is the logic of the "think first" step and of the distinction between scaffold and crutch discussed in Section 4.4. Related arguments about learning through the child's own observation, activity and discovery were made by Pestalozzi (1894), Montessori (1912), Bruner (1961), and Freire (1970).

These educators did not know AI, but their principles remain relevant to the questions it raises. Good teaching rests on questions rather than ready-made answers, on knowledge of each pupil, on trust between teacher and child, and on a teacher who knows when to help and when to step back. Current AI systems cannot replace the relational, contextual and ethical judgement

of the teacher. They can, however, become a valuable assistant to the teacher who exercises it. The protocol proposed in the next section applies these principles to a new tool.

6 A proposal for the classroom and the school

6.1 The "think – ask – check" protocol

Based on the above, a simple three-step protocol is proposed for every activity involving AI. It is a proposal of this paper: it has not been tested as a whole, and its effects have not been evaluated. The protocol connects the three issues of this paper:

(1) Think first on your own. Before any question is put to the AI, the pupil formulates their own answer, hypothesis or first draft, however imperfect. This preserves the "desirable difficulty", and AI comes in to be compared with the child's thinking, not to replace it.

(2) Ask wisely. The question is phrased without personal information, in line with the class privacy card (Section 3.3). In primary school, the question is typed by the teacher, while its wording is decided jointly by the class.

(3) Check before you trust. The answer is compared with an independent source, such as the school textbook, a second text or the child's own first answer. The pupil explains what they kept, what they rejected and why.

Assessment of the activity focuses on the third step: what is graded is not the final product but the pupil's justification of what they accepted or rejected from the AI.

6.1.1 How the protocol was derived

The protocol was not designed in advance as a research instrument; it is the author's synthesis of elements of the literature reviewed above and of her own practice. It was not developed through formal cycles of testing and revision; it took shape afterwards, from the author's reflection on what had interested the pupils and how they had responded in the lessons. From the literature, the "think" step draws on research on desirable difficulties (Bjork, 1994) and on the finding that AI offering hints rather than solutions did not harm exam performance (Bastani et al., 2025); the "ask" step applies the principle of data minimisation (Regulation, 2016) and the Greek ministerial guidance (Ministerial Decision, 2026) to the way questions are phrased; and the "check" step draws on research on automation bias (Parasuraman & Manzey, 2010) and on definitions of critical thinking that emphasise justification (Dewey, 1910; Facione, 1990). From practice, the "check" step and the assessment of justification were used in the lessons described in Section 2.2 (comparison with the school textbook and the water-shortage scenario), while the "ask" step corresponds to the data-protection activities of the lesson scenario described in Section 3.3. The "think first" step, and the protocol as a complete sequence, are proposed here and have not been implemented or evaluated as such.

6.1.2 What the protocol adds

Its individual elements are not new. AI-literacy frameworks (Long & Magerko, 2020; Miao & Shiohira, 2024) already emphasise the critical evaluation of AI outputs; guided discovery (Bruner, 1961) emphasises the learner's own effort before receiving an answer; and scaffolding (Wood et al., 1976) emphasises temporary support that is gradually withdrawn. The contribution of the protocol is modest and practical. It (a) combines cognitive and data-protection concerns in a single routine, so that privacy-conscious questioning becomes part of every AI activity rather than a separate lesson; (b) translates these principles into three short steps that primary teachers can use in an ordinary lesson in which the AI is operated by the teacher, consistent with UNESCO's recommended minimum age for independent use (Miao & Holmes, 2023); and (c) moves assessment from the AI-assisted product to the pupil's justification. Whether this combination is effective is an empirical question that remains open. Table 1 gives a proposed adaptation of the protocol by age. It is not based on evidence of an age progression: the grade bands follow the structure of the Greek primary curriculum and are intended as a starting point for teachers and for future research.

6.2 The role of school leadership

School leadership can help turn the principles set by the State into everyday school practice. By decision of the Teachers' Council, the school can agree on short common rules for the use of AI, inspired by the GDPR and the principles of Ministerial Decision 55909/Δ6/2026 and adapted to the age of primary school pupils. Indicatively, these rules may include the following. No one, including teachers, enters pupils' identifying data into AI tools. In the classroom, AI

Table 1 Proposed (untested) adaptation of the "think – ask – check" protocol by age

Grades (age)	Think	Ask	Check
Grades 1–2 (age 6–8)	The child says orally what they believe or draws their idea	The teacher asks the AI in front of the class; the children suggest the question	Comparison with a picture, a book or the children's experience: "Does this agree with what we know?"
Grades 3–4 (age 8–10)	A short written prediction or answer before using the AI	The class words the question together and removes any personal information	Comparison with the school textbook; recording one difference and one similarity
Grades 5–6 (age 10–12)	A first draft of a solution or text by the pupil	Groups draft and refine the question; privacy card	Cross-checking with two sources; written justification of what was accepted or rejected

Notes: Proposal of the present paper; it has not been empirically tested and does not describe an evidence-based developmental progression. In all grades, the AI is operated by the teacher. Greek primary school comprises six grades (Grades 1–6) for children aged 6–12.

is used only through the teacher. Photographs or recordings of pupils are not posted without written consent, and only the minimum necessary material is ever posted. Parents are informed about the rules and about the limits of AI use at home.

School leadership also has a role in informing parents. Many parents themselves need information in order to support their children at home. A short parents' meeting on "What AI does and does not do" can carry the "think – ask – check" protocol into the home.

7 Discussion

The analysis suggests that the three issues addressed in this paper are aspects of the same pedagogical need: for the pupil to remain an active subject in the face of a tool that can easily turn them into a passive recipient. A pupil who questions an answer is exercising critical thinking. When they consider what they should not reveal, they are protecting their data. When they answer on their own first, they are protecting their thinking. The paper proposes that these three attitudes can be addressed together and with the same means: guided discovery, an independent source for comparison, and justification. The author's classroom experience illustrates that such activities can be implemented in ordinary primary school lessons, but their effects on pupils' critical thinking have not yet been evaluated.

The fear of cognitive weakening should not be underestimated, but neither should it turn into moral panic. The history of educational technology suggests that many new tools provoked similar concerns and that the answer was usually found not in prohibition but in the pedagogical design of its use. The study by Bastani et al. (2025) conveys a similar message: the same language model, under different rules of use, had different effects on the exam performance of high school students. It is the teacher who sets these rules in the classroom. As in the car example, an excellent tool does not in itself guarantee a safe journey; responsibility lies with whoever holds the wheel.

The paper has limitations. The review is narrative and not systematic, and the selection of sources reflects the author's judgement; research on generative AI is evolving rapidly, and most empirical data concern adults or adolescents, often through self-report and over short periods. The classroom examples are the author's practitioner reflections from a single school; they were not collected as research data and are used for illustrative purposes only. The "think – ask – check" protocol and its age adaptation have not been tested. Future research should examine, with appropriate ethical approval and informed parental consent, whether the protocol, and guided use of AI more generally, affects primary school pupils' task performance, learning retention, critical-thinking disposition and cognitive offloading, treating these as distinct outcomes and using larger samples and repeated measurements.

8 Conclusion

AI is already in school, and primary school children encounter it every day. The question is not whether they will use it, but whether they will learn to use it with judgement. The author's classroom experience illustrates that activities which encourage pupils to question AI answers can be implemented in ordinary primary school lessons, even in the early grades, with simple means such as comparison with the school textbook, contrasting the AI's answer with the children's own experiences, and assessing justification rather than the final product. Their effects on pupils' critical thinking, however, have not yet been evaluated. The regulatory

framework provides clear principles for protecting minors' data, which can be translated into rules that children understand and into professional habits for teachers. Finally, research on cognitive dependence, which concerns mainly adolescents and adults, suggests that AI that gives ready-made answers may pose a greater risk to learning than AI that guides; whether this holds for primary school pupils remains to be investigated.

The "think – ask – check" protocol sums up the position of this paper, and is offered as a proposal for teachers to try and for researchers to test: first the child thinks, then asks the machine, and finally judges for themselves. The rules and tools of the State set the framework. How teaching takes place within it remains, as always, a matter of the teacher's pedagogical judgement. In this, long-standing educational principles remain a reliable guide. AI can become a wonderful assistant in this work, provided that control remains in the teacher's hands and thinking remains in the child's mind.

Ethics statement

This paper is a perspective piece and does not report a research study. The classroom activities described in it arose in the course of the author's ordinary teaching and were not planned as research: they were lessons within the official curriculum, carried out with whole classes, whose only aim was to teach the pupils to use AI critically and to protect themselves when using digital tools. The decision to write about these lessons was taken later, when the author decided to share with colleagues what she had learned from them. Because no research was planned, no ethics approval was sought and no institutional authority or ethics body was asked to determine whether approval was required; for the same reason, parents and pupils were not informed about, or asked to consent to, any later use of classroom material. Accordingly, the paper reports no data derived from pupils. No notes were kept during the lessons and no material produced by pupils was collected, since no later use was intended. The descriptions are based on personal notes that the author wrote from memory afterwards, when she decided to summarise her experience, and on her later reflection, as is usual in practitioner accounts; these notes are the author's own reflections and are not used as research data. The approximate figures, the table of results, the class sizes and the remarks on parents' occupations that appeared in the original submission have been removed. The lessons are described only qualitatively, at the level of the class and from the author's perspective as the teacher; pupils' words are paraphrased rather than quoted, and no individual pupil or individual response can be identified.

Declaration on the use of AI tools

During the preparation of this paper, the tool Claude (Anthropic) was used to draft a first version of the text based on the author's own material and reflections, to search for and check the literature, to format references, and to translate the manuscript from Greek into English; it was also used to help prepare the revision in response to the reviewers' comments. The entire text was reviewed and revised by the author, who takes full responsibility for the accuracy, validity and originality of its content.

Conflicts of interest

The author declares no conflict of interest.

References

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). https://doi.org/10.1073/pnas.2422633122

Bjork, R. A. (1994). Memory and Metamemory Considerations in the Training of Human Beings. Metacognition, 185–206. https://doi.org/10.7551/mitpress/4561.003.0011

Bruner, J. S. (1961). The act of discovery. Harvard Educational Review, 31(1), 21–32.

Dahmani, L., & Bohbot, V. D. (2020). Habitual use of GPS negatively impacts spatial memory during self-guided navigation. Scientific Reports, 10(1). https://doi.org/10.1038/s41598-020-62877-0

Dewey, J. (1910). How we think. (, Ed.). D C Heath. https://doi.org/10.1037/10903-000

Facione, P. A. (1990). Critical thinking: A statement of expert consensus for purposes of educational assessment and instruction (The Delphi Report). The California Academic Press.

Freire, P. (1970). Pedagogy of the oppressed (M. B. Ramos, Trans.). Herder and Herder.

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

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of Hallucination in Natural Language Generation. ACM Computing Surveys, 55(12), 1–38. https://doi.org/10.1145/3571730

Kareli, F. (2026, November 20–22). Artificial intelligence as an assistive technology for the inclusion of students with learning difficulties and ADHD: Eleven examples of application in primary school [Unpublished conference paper, accepted for oral presentation; proceedings forthcoming]. 49th International Conference of Humanities and Social Sciences "Special Needs Sciences", Institute of Culture, Democracy and Education, Patras, Greece (online). [In Greek]

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274. https://doi.org/10.1016/j.lindif.2023.102274

Katsampoxaki-Hodgetts, K., Kotsidis, K., Papadakis, S., & Anastasiades, P. (2025). Reimagining preschool pedagogy: AI-enhanced teaching through multiliteracies-informed curriculum innovation. In V. Lamanauskas (Ed.), Science and technology education: Expectations and experiences (Proceedings of the 6th International Baltic Symposium on Science and Technology Education, BalticSTE2025) (pp. 104–119). Scientia Socialis. https://doi.org/10.33225/BalticSTE/2025.104

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 [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2506.08872

Law. (2019). 4624/2019. Hellenic Data Protection Authority, measures for implementing Regulation (EU) 2016/679 and transposition of Directive (EU) 2016/680 into national law. Government Gazette of the Hellenic Republic, A 137, 29.08.2019. [In Greek]

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 (pp. 1-22). https://doi.org/10.1145/3706598.3713778

Livingstone, S., Stoilova, M., & Nandagiri, R. (2019). Children's data and privacy online: Growing up in a digital age. An evidence review. London School of Economics and Political Science.

Long, D., & Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 1–16. https://doi.org/10.1145/3313831.3376727

Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. https://doi.org/10.54675/ewzm9535

Miao, F., & Shiohira, K. (2024). AI competency framework for students. UNESCO Publishing. https://doi.org/10.54675/jkjb9835

Ministerial Decision. (2026). 55909/Δ6/2026. Determination of issues relating to the use of artificial intelligence systems in the educational process and the protection of personal data. Government Gazette of the Hellenic Republic, B 2639, 12.05.2026. [In Greek]

Montessori, M. (1912). The Montessori method (A. E. George, Trans.). Frederick A. Stokes.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and Bias in Human Use of Automation: An Attentional Integration. Human Factors: The Journal of the Human Factors and Ergonomics Society, 52(3), 381–410. https://doi.org/10.1177/0018720810376055

Pestalozzi, J. H. (1894). How Gertrude teaches her children (L. E. Holland & F. C. Turner, Trans.). Swan Sonnenschein.

Regulation (EU). (2016). 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union, L 119, 4.5.2016.

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L, 2024/1689, 12.7.2024.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive Offloading. Trends in Cognitive Sciences, 20(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. Science, 333(6043), 776–778. https://doi.org/10.1126/science.1207745

Uğraş, H., Uğraş, M., Papadakis, S., & Kalogiannakis, M. (2024). Innovative Early Childhood STEM Education with ChatGPT: Teacher Perspectives. Technology, Knowledge and Learning, 30(2), 809–831. https://doi.org/10.1007/s10758-024-09804-8

UNICEF. (2021). Policy guidance on AI for children (Version 2.0). UNICEF Office of Global Insight and Policy.

Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Harvard University Press.

Wood, D., Bruner, J. S., & Ross, G. (1976). The Role of Tutoring in Problem Solving. Journal of Child Psychology and Psychiatry, 17(2), 89–100. Portico. https://doi.org/10.1111/j.1469-7610.1976.tb00381.x

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(1). https://doi.org/10.1186/s40561-024-00316-7