Pedagogika, 2025, 75(4), 303–308
http://pages.pedf.cuni.cz/pedagogika/

Artificial Intelligence, Education, and the Need for Knowledge Consilience in the Post-digital Era

Martin Malčík, Tomáš Zemčík

Artificial intelligence is no longer a distant technological promise or an exclusively engineering problem. It has become a structural force that increasingly influences decision-making processes, redistributes power over information, and transforms the institutional foundations of contemporary societies. Education – as a social institution and cultural practice – stands at the very centre of this transformation. This special issue stems from the conviction that the relationship between artificial intelligence and education can no longer be meaningfully grasped within individual disciplines or through established methodological traditions.

The basic premise of this issue is the idea of the consilience of knowledge, which was formulated by William Whewell in the 19th century and significantly revived by Edward O. Wilson in the 20th century. Consilience refers to the convergence and interconnection of knowledge across disciplines – a process in which different epistemic perspectives come together to provide a better understanding of complex phenomena that cannot be adequately explained in isolation within the boundaries of specific scientific disciplines and their sub-disciplines. Artificial intelligence, especially in its educational, social, and ethical dimensions, is precisely such a phenomenon.

Educational sciences traditionally belong to the humanities and social sciences. They are grounded in interpretive frameworks, normative thinking, and an emphasis on meaning, understanding,

DOI: 10.14712/23362189.2026.5207

and human experience. In contrast, artificial intelligence arises from the technical sciences based on formal logic, statistical learning, and optimisation models. It is becoming increasingly clear, however, that these two areas can no longer be kept separate. Artificial intelligence systems are now directly involved in educational processes, teaching environments, assessment, material preparation, institutional management, and decisions about individuals' educational paths, not to mention students' use of artificial intelligence as a tool for doing their homework. At the same time, education fundamentally shapes how societies understand artificial intelligence, how they regulate it, and how they are able to live with it.

As Yuval Noah Harari convincingly demonstrates in his book *Nexus*, artificial intelligence systems already participate in decisions that have a direct and profound impact on human lives. In the United States and elsewhere, algorithmic systems are used in criminal law (e.g. the COMPAS tool), in decisions concerning admission to universities or employment, in assessing eligibility for social benefits, in approving bank loans, in medical diagnoses, and in mediating romantic relationships through platforms such as Tinder. These examples illustrate a broader process: traditional institutions that have historically managed information and decision-making mechanisms are disintegrating or being replaced by algorithmic information networks.

Algorithmic information networks are not primarily designed as tools for producing truth, facts, or verified information, but rather as systems that generate the illusion of credible outputs. Their functional principle is not epistemic validation, but the probabilistic generation of language and data patterns that may be correct, partially misleading, counterfactual, or completely false. This phenomenon is often referred to as "hallucination" in the context of generative artificial intelligence; however, it is not an exception or a system failure, but a structural feature of such systems. The decisive factor is therefore not the existence of the hallucination itself, but rather the human attitude towards it: if the output that is generated corresponds to the user's expectations, it is interpreted as true; if not, it is retrospectively labelled as incorrect or hallucinatory.

This reality poses a fundamental challenge for professions in which the consequences of erroneous conclusions are socially significant – for lawyers working with algorithmically generated references to legal norms and case law and doctors using AI to formulate diagnoses and prognoses, but also for educators, for whom artificial intelligence is increasingly helping to shape teaching content and materials. The responsibility for verifying sources, the validity of information, and the adequacy of conclusions does not lie with the technology itself, but with the expert who uses AI as a tool.

This places increased demands on the development of critical assessment skills, rapid and thorough verification of information, subject expertise, and professional judgement – skills that cannot be automated and which are becoming a key prerequisite for the meaningful and responsible use of artificial intelligence in a post-digital or post-factual society.

In this sense, critical assessment and systematic verification should not be understood merely as one set of skills among many required in the age of artificial intelligence. Rather, they constitute a central epistemic condition for any meaningful engagement with AI in education. Without such critical capacities, neither the consilience of knowledge nor the development of new interpretative frameworks across disciplines is possible. Critical thinking thus emerges not only as an educational objective, but also as a prerequisite for responsible participation in a post-digital society increasingly shaped by algorithmic decision-making.

This creates an institutional and epistemic interspace in which decision-making power is shifting faster than we are able to develop the corresponding interpretative, legislative, regulatory, and ethical frameworks. Technological experts who design neural networks and machine learning models are undoubtedly crucial. However, they cannot alone provide answers to questions about the meaning, values, responsibilities, and impacts of these systems on human development, education, democracy, and social cohesion, let alone offer examples of good practice. At the same time, researchers in the humanities and social sciences, including educators and teachers, increasingly find themselves in a situation where their professional practice is shaped by technologies whose internal logic they do not fully understand.

It is precisely in this interspace that an increased risk of errors, simplifications, and inadequate conclusions arises. Non-experts in the field of artificial intelligence are forced to work with systems that fundamentally affect their everyday practice without having sufficient tools for their critical evaluation. At the same time, the pace of technological development does not allow us to wait until methodological approaches stabilise and "good practice" emerges. Waiting for consensus in this situation has become a luxury we can no longer afford.

This situation calls for the development of new interpretative frameworks – frameworks that are not merely a simplified translation of technical concepts into educational language, but which focus critically on how artificial intelligence is transforming epistemic authority, the educational environment, and the very conditions of education. These frameworks can only emerge in dialogue across disciplines, in the confrontation of empirical knowledge with philosophical reflection, ethical analysis, and educational reality, theory, and practice.

This special issue is conceived as a response to this urgency. It creates a space for experts in educational sciences and other related fields to formulate rigorous yet open and exploratory perspectives on artificial intelligence in education and on education about artificial intelligence. Importantly, this issue does not assume that authors must be technical specialists in AI. On the contrary, it recognises the legitimacy and necessity of voices that reflect on artificial intelligence from the perspective of their own professional expertise.

The contributions collected in this issue show that consilience is not a finished methodological programme, but an ongoing process. There is as yet no single way to scientifically grasp the full extent of the impact of artificial intelligence on education. Methods, concepts, and research designs are still taking shape. However, this absence of methodological closure is not a weakness, but a reflection of the novelty and complexity of the phenomenon itself.

Education has always been a space in which societies articulate their hopes, fears, and normative ideas about the future. Artificial intelligence enters this space not merely as a tool, but as a co-creator of educational reality. If educational research is to remain socially relevant, it must relate to artificial intelligence not only instrumentally, but also critically and reflectively. It must help define what is meaningful, how far thinking can go, and where new boundaries need to be set.

This special issue therefore offers a collective intellectual intervention at a moment when the rules of the game are not only changing but have already changed! Artificial intelligence is not coming – it is already here. The editors are aware that the contributions in this issue will not provide definitive answers. Rather, they aim to foster an interdisciplinary dialogue in which there is room not only for traditional scientific conclusions, but also for conclusions and exaggerated interpretations that may prove inadequate, misguided, or erroneous over time. It is precisely this kind of "experimental" thinking that helps to map the horizon of meaning to the limits of nonsense. Without such reflection, it is impossible to respond adequately to the challenges associated with artificial intelligence in education.

What does this issue say about the state of artificial intelligence in education?

The opening article by Tomáš Zemčík addresses the issue of defining artificial intelligence (AI) in a historical, philosophical, technological, and legislative context and shows that there is not – and cannot be – a single universally valid definition of AI. The author traces the development of definitions from the early optimistic ideas of strong AI, through periods of so-called "AI summers and winters", to the current era of data-driven models and deep learning, emphasising the so-called *AI effect*, which shows that technology ceases to be perceived as AI once it becomes a common part of practice. Critically, the text distinguishes between human intelligence and AI as an artificial imitation of intelligence and points out the illusory nature of the "understanding" of the ethical neutrality of algorithmic outputs.

The article by Lukáš Honzík et al. examines the limitations of generative artificial intelligence, specifically the ChatGPT language model, in solving classic and modified logic tasks such as "The Wolf, the Goat, and the Cabbage" or "Cannibals and Missionaries". The authors show that while AI can offer correct solutions to well-known and frequently-described tasks, it repeatedly makes errors of logic when these tasks are modified or made more complex, fails to recognise them reliably, and sometimes

produces what are called hallucinations. These limitations are explained, among other things, by the Chinese room argument, which illustrates the difference between truly understanding a problem and merely following rules formally.

The next article, by Michal Místecký, examines the extent to which GPT language models trained on the works of specific authors are able to truly imitate their literary style, using methods of quantitative linguistics and stylometry. Drawing on a large corpus of texts by four Czech authors (Čapek, Hašek, Kafka, Vančura), the author compares original literary works with texts generated by corresponding "authorial" GPT models through morphological, syntactic, and lexical indicators and analyses of the most frequent linguistic elements. The results show that although the texts generated by the models resemble those of the authors in some respects, they are in most cases statistically and stylometrically easily distinguishable from the human originals.

The article by Miroslava Černochová et al. analyses national policies for the integration of generative artificial intelligence (GenAI) into primary and secondary education in four countries: the Czech Republic, Israel, Latvia, and Slovakia. The study identifies significant differences in the level of preparedness of the individual countries. Israel stands out with its centralised approach and comprehensive national strategy covering the curriculum, ethical rules, and teacher training. Latvia is integrating GenAI into a broader digital transformation as part of *the Skola2030 reform*, while the Czech Republic and Slovakia still lack comprehensive national strategies and rely primarily on initiatives by non-profit organisations and universities. The authors emphasise that the successful implementation of GenAI requires not only technological readiness, but, above all, coherent policies, high-quality teacher training, ethical standards, and equal access to digital infrastructure.

The study by Tomáš Botlík Nuc and Marie Klusáčková examines the impact of AI-generated songs on vocabulary acquisition and student motivation in German as a second foreign language at a secondary vocational school in Prague. In a quasi-experimental study with 60 observations, the traditional way of teaching vocabulary was compared with an innovative approach using the Suno AI platform, which allows students to create their own songs with the target vocabulary. The results showed a statistically significant difference in favour of the experimental group, which achieved an average score of 6.2 points compared to 4.2 points in the control group.

The study by Klára Rybenská and Sylva Sklenářová discusses the use of generative artificial intelligence in the digitisation and transcription of printed and handwritten historical documents, with a focus on OCR (Optical Character Recognition) and HTR (Handwritten Text Recognition) technologies. The research tested three AI tools – ChatGPT-4, Claude Sonnet, and Copilot – on seven handwritten texts of varying complexity dating from 1980 to 1990, with Claude achieving the best results in transcription and analysis accuracy, ChatGPT providing high-quality interpretations

with lower transcription accuracy, and Copilot lagging significantly behind in all categories. The research highlights the educational potential of artificial intelligence for making historical documents accessible to the general public and students, but also draws attention to ethical aspects, in particular the need to preserve the authenticity of historical data and critically verify AI outputs.

In addition to six peer-reviewed articles, the special issue also includes four non-peer-reviewed texts, all of which focus on various practical aspects of the use of AI in education.

Finally, we would like to thank all the authors who, in a relatively short time after the call for papers was published, provided high-quality texts reflecting the themes of the special issue. We would also like to express our gratitude to the reviewers, whose thoughtful and rigorous evaluations helped to make the texts more accurate and concise. We hope that readers will gain as much inspiration as possible, both for a theoretical understanding of issues related to artificial intelligence and for practical ideas for its use in the school environment.

Martin Malčík
VSB – Technical University of Ostrava, Department of Social Sciences;
email: martin.malcik@vsb.cz

Tomáš Zemčík
VSB – Technical University of Ostrava, Department of Social Sciences;
email: tomas.zemcik@vsb.cz