Cognitive Symbiosis or Cognitive Dependence: The Inherent Tensions and Governance Pathways of Generative Artificial Intelligence in Reshaping Education

Cao Hongyu[1], Gan Weixiang[2]*

1. Zhenguigu Technology Co., Ltd. No. 55, Langfang, Hebei 065000, China

2. Lincoln University College, Petaling Jaya, Selangor Darul Ehsan-47301, Malaysia

*Email: Gabriel1995@qq.com

Abstract: Generative artificial intelligence is profoundly reshaping the cognitive ecology of education. Yet much mainstream research remains confined to a framework of "tool-based augmentation", treating it as an external means of improving efficiency and therefore struggling to explain deeper changes in human-machine relations. This paper argues for moving beyond instrumentalism: generative AI should be reconceived not as a teaching tool external to the learner, but as a "cognitive mediator" that participates in the learner's construction of meaning. It then uses the paired categories of "cognitive symbiosis" and "cognitive dependence" to characterise two possible trajectories of human-machine collaborative learning. The analysis shows that the educational use of generative AI contains three interrelated tensions: cognitive augmentation versus degradation at the individual level, efficiency gains versus widening divides at the social level, and personalisation versus homogenisation at the level of knowledge. These tensions share a common source: when machine output replaces human cognitive labour without preserving human agency, symbiosis degenerates into dependence. On this basis, the paper proposes "responsible cognitive symbiosis" as a governance orientation, to be pursued through coordinated action in AI literacy, instructional and assessment design, and institutional safeguards. It concludes that the fundamental task of education in the age of AI is not to compete with machines, but to protect those human capacities that cannot be replaced by algorithms, including judgement, meaning-making, and care.

Keywords: generative artificial intelligence; cognitive symbiosis; cognitive dependence; digital transformation of education; AI literacy; human-centred AI

0. Introduction

Since generative artificial intelligence represented by ChatGPT entered public view at the end of 2022, it has permeated the everyday practices of education with remarkable speed. Approximately two months after launch, ChatGPT had already surpassed 100 million monthly active users[1]. Surveys indicate that a substantial proportion of university students now use such tools in their studies[2]. In parallel, academic research on the educational applications of AI has grown rapidly in recent years[3][4], and Chinese scholars have also engaged intensively with its opportunities, challenges, ethical implications, and governance pathways[5]. Education is among the first fields to be affected because its core tasks, namely the transmission of knowledge and the cultivation of minds, fall squarely within the capability range of generative AI. Yet policy documents and most empirical studies have tended to respond to this disruption through an "instrumental augmentation" logic, viewing AI as a means to improve efficiency, realise personalised instruction, and reduce teachers' workload[6][7]. This approach has its value, but it rests on an unexamined premise: that AI always remains external to the learning subject, and that its intervention does not change the nature of learning and cognition. Critical research on educational technology has long questioned this presumption of "technological neutrality"[8][9], while assessments of AI's educational impact remain sharply divided between optimistic narratives of liberation and pessimistic narratives of replacement[10].

In China, this issue is also intertwined with the national strategy of educational digitalisation and the broader project of building a strong education system. At the policy level, high expectations have been placed on AI-enabled education, which is seen as an important lever for improving both educational equity and quality. At the same time, the rapid spread of generative AI has made questions of academic integrity, data security, and the protection of minors newly urgent[11]. In this context, neither simple prohibition nor laissez-faire use is adequate. Prohibition is unrealistic and may drive students to use these tools privately without guidance; permissive use, meanwhile, may quietly erode education´s most fundamental formative function under the banner of efficiency. How to strike a balance between embracing technology and safeguarding education´s human-forming purposes has therefore become a genuine question that cannot be avoided or answered lightly. This dilemma is not unique to China but is shared by education systems across the world[12]. Nevertheless, China´s vast learner population and the developmental disparities between urban and rural areas and across regions amplify the consequences of relevant decisions. Clarifying the theoretical relationship between AI and education is therefore not a purely academic exercise: it bears directly on the kind of cognitive environment in which hundreds of millions of learners will grow.

The crux lies precisely here. Generative AI differs from earlier educational technologies. It is not a device that passively executes predefined tasks; it can actively generate content and intervene conversationally in the entire process through which learners think. Once learners form a durable cognitive coupling with such a machine, how to use it to improve teaching and learning. A more fundamental question must be asked: how will human cognition be reshaped after being embedded over time in an environment constituted by generative machines? Will it move towards the collaborative enhancement of capability, or towards dependence on the machine and the withdrawal of human judgement? This paper develops its argument through the paired categories of cognitive symbiosis and cognitive dependence. It first analyses the paradigm shift triggered by the entry of AI into education, then examines the cognitive mechanisms and conditions of human-machine collaborative learning, before revealing three latent tensions and discussing possible governance pathways. The paper does not seek to choose between optimism and pessimism. Rather, it argues that symbiosis or dependence is not determined by technology alone, but by the conditions under which human-machine interaction occurs. It can therefore be shaped by education and governance.

1. From Tool to Cognitive Mediator: A Paradigm Shift in the Relationship Between AI and Education

Before the emergence of generative AI, AI in education mainly appeared in the form of specialised applications such as intelligent tutoring systems, adaptive learning systems, and learning analytics tools. Their basic logic was to collect learner data, use it to make predictions, and then implement interventions[9]. Luckin, Holmes, Griffiths, and Forcier[13] once summarised three metaphors for the role of AI in education: tutor to the student, tool for learning, and object taught by the student. Although these metaphors differ, they share the same underlying assumption: AI is an external device with clear boundaries and determinate functions, serving pre-set human goals, and its operation is explainable, predictable, and controllable. This instrumental view is deeply rooted in the tradition of educational technology research. It presumes that human beings are the subjects who use technology, while technology is the object being used[8]. From this perspective, the measure of AI contracts into a single dimension: whether it enables an established task to be completed faster, more cheaply, and more accurately. It should be acknowledged that this judgement is broadly valid for specialised systems such as adaptive tutoring, and such systems have indeed shown effects. The real question is whether the same judgement can be applied unchanged to a technology of a fundamentally different kind.

Even during the heyday of instrumentalism, the actual effects of AI applications in education

fell far short of their promises. On the one hand, intelligent tutoring systems have proven effective in well-structured subjects, and adaptive practice can to some extent dynamically match task difficulty to learner performance. On the other hand, many products have remained at the demonstration or pilot stage for long periods, lacking rigorous evidence of impact, and a significant gap has persisted between research and practice[3][4]. This gap itself is a warning: simply equating AI with a tool for improving efficiency both overestimates technology and narrows education. More importantly, instrumentalism cannot explain a newly emerging fact: learners are already using generative AI, on a large scale and often spontaneously, in the core stages of learning rather than merely for peripheral auxiliary tasks. The arrival of generative AI has not ended the debate over technology and education. It has reopened that debate in a sharper form, because what it touches is no longer the efficiency of one teaching procedure but the basic mode of learning and cognition itself.

Several features of generative AI are loosening the assumptions of the instrumental view. It has generality: it was not designed for one specific task, yet it can generate content across almost all domains. It has generativity: it does not merely retrieve and present existing information, but participates directly in the production of meaning. It has dialogicity: through natural language, it becomes continuously embedded in the unfolding of thought. To understand this shift, the term "tool" has become insufficient; a more explanatory perspective is needed. Vygotsky´s[14] theory of mediation shows that higher psychological functions are realised through mediators such as signs and tools. Tools do not merely change the efficiency of action; they reshape the structure of cognitive activity. The theories of distributed cognition[15][16] and the extended mind[17] further argue that when external devices participate stably in cognition, they become functional components of the mind. This insight is not new. More than three decades ago, Salomon, Perkins, and Globerson[18], in their discussion of "cognitive partners", distinguished between the effects produced with technology and the effects retained after technology is withdrawn, reminding us that the educational value of intelligent technologies lies not in completing tasks for humans but in how they change the minds of their collaborators. Following this line of thought, generative AI is better understood as a generative cognitive mediator rather than as yet another convenient tool.

Viewing AI as a cognitive mediator has two consequences. First, the unit of analysis shifts from the isolated individual learner to the coupled learner-AI system. The evaluative question also changes: it is no longer enough to ask what AI does for learning; we must also ask what it does to the learner´s mind. Second, the conception of knowledge becomes unsettled. In the traditional picture, "knowing" means that knowledge is internalised through learning as an internal structure that can be called upon when needed. In an environment built by generative AI, however, "knowing" is quietly sliding towards "knowing how to retrieve". Knowledge seems no longer to require internalisation; it can simply be called up instantly. This shift is not wholly negative. Using external resources is a normal condition of human cognition[18]. Its potential cost is that if internalisation is repeatedly bypassed, learners will find it more difficult to weave scattered information into coherent understanding and more difficult to think independently without external support. For this reason, Heersmink[19] argues that education in relation to the Internet and intelligent tools should cultivate the relevant intellectual virtues, enabling people to use tools well without becoming passive recipients of information. This paper refers to the desirable state that human-machine cognitive coupling may achieve as cognitive symbiosis, and this concept guides the discussion that follows.

The perspective of cognitive mediation also requires us to rethink the position of the teacher. Under instrumentalism, the teacher is the user of technology and the operator of the classroom. In the picture of human-machine cognitive coupling, the teacher is better understood as the designer and guardian of the coupled system. The teacher must decide where AI should be introduced and where deliberate space should be left; must model how to converse with machines and how to maintain the necessary vigilance towards their outputs. From this perspective, the more AI can

take over the transmission of information, the more the teacher´s irreplaceability becomes visible in those areas beyond the machine´s reach: value guidance, emotional support, and sensitive attention to the specific difficulties of particular students[20]. This also suggests that the popular question of whether AI will replace teachers may have been misdirected from the start. The real question is not whether replacement will occur, but how the professional role of teachers will be redefined in the division of labour with AI. Unfortunately, existing research has focused more on technologies and systems than on the position of teachers and pedagogy[3], even though this is precisely what determines whether technology can genuinely serve educational formation.

2. Why Cognitive Symbiosis Is Possible: Mechanisms and Conditions of Human-Machine Collaborative Learning

Cognitive symbiosis refers to a condition in which AI undertakes part of the cognitive load and lower-order operations, thereby allowing human cognitive resources to be directed towards higher-order understanding, judgement, and creativity, while humans retain metacognitive monitoring over the cognitive process. It is adjacent to, but not identical with, cognitive offloading. Offloading means transferring cognitive tasks that one would otherwise perform oneself to external resources, such as using notes instead of memory or a calculator instead of mental arithmetic[21]. Its value is originally neutral. If the resources released by offloading are reinvested in more valuable thinking, offloading can lead to symbiosis; if it merely produces idleness and atrophy in the corresponding capability, it turns into dependence. Symbiosis and dependence are therefore not two different technologies, but two possible outcomes of the same human-machine interaction under different conditions. This also means that locating either hope or fear in technology itself is a mistake. The more important question is which conditions can reliably guide human-machine interaction towards symbiosis. Several classical theories in educational psychology provide clues for answering this question, while also revealing the demanding preconditions under which symbiosis can be achieved.

To clarify the distinction between symbiosis and dependence, offloading itself must be analysed more carefully. Transferring tedious operations that are not directly related to learning objectives, such as formatting, information retrieval, or grammar correction, is usually beneficial, because it releases limited cognitive resources from low-value labour. The problem arises with another kind of offloading: when learners also hand over core cognitive tasks that should be completed by themselves, such as understanding, reasoning, argumentation, and judgement, what is removed is no longer a burden but learning itself. The first type of offloading serves learning; the second replaces learning. They are similar in form but different in substance. The real difficulty is that the convenience of generative AI blurs this boundary. It makes substitutive offloading so easy and tempting that learners often cross the line without noticing. This explains why the same technology can support symbiosis or slide into dependence: the difference never lies simply in whether offloading is used, but in which cognitive activities are offloaded and whether the learner remains clearly aware of this distinction.

Consider first the zone of proximal development and scaffolding. Effective teaching provides appropriate support between the learner´s independent level and potential level. The promise of generative AI lies in its possibility of providing personalised scaffolding for large numbers of learners at low marginal cost, thereby approaching Bloom´s[22] "two-sigma" ideal: learners receiving one-to-one tutoring could perform around two standard deviations better than those taught in conventional classes. This goal, long regarded as difficult to scale, appears for the first time to have a technological basis. Yet scaffolding theory also requires that support be gradually withdrawn as competence develops. If scaffolding is never removed, it is transformed from assistance into a prop. Cognitive load theory and desirable difficulty point in the same direction. Sweller[23] distinguished between extraneous load that hinders learning and load germane to schema construction, while Bjork and Bjork[24] showed that retrieval practice, spacing, moderate

confusion, and productive struggle may make learning effortful but are also conditions for long-term memory and deep understanding. The paradox of generative AI is that it is so skilled at smoothing the learner's path that, while reducing extraneous load, it often also removes the desirable difficulties that should have remained.

Imagine two students facing the same essay-writing task. One first develops an outline and preliminary argument independently, then asks AI to challenge the weak points, supply counterexamples, and support repeated revision. In this process, the machine plays the role of an opponent who pushes the student to think more deeply. The student's judgement is not replaced but exercised. The other student directly asks the machine to generate the entire text, submits it after minor wording adjustments, and receives an apparently acceptable assignment without having undergone any substantive thinking. The two students have used the same tool, but their cognitive gains are entirely different. This contrast matters because it shows clearly that what determines the success or failure of learning is not whether AI is used, but how it is used and who remains the subject of thinking during its use.

The pivot determining the direction of development is self-regulated learning. Metacognition, that is, awareness and regulation of one's own cognition[25], enables learners to set goals, monitor processes, and adjust strategies accordingly[26]. In human-machine collaboration, this capacity determines what role the machine plays. Learners with stronger self-regulation tend to treat AI as an interlocutor who questions and refutes them, thereby deepening rather than replacing thought. Mollick[27] calls this mutually enhancing collaboration "co-intelligence". Learners with weaker self-regulation, by contrast, are more likely to treat it as an all-answering oracle and hand over the whole labour of thinking. The same conversational interface can function either as a Socratic questioner or as an answer machine that does everything on the learner's behalf. The dividing line lies not in the technology itself but in the user's intention and capacity. It should also be acknowledged that empirical evidence on the long-term effects of generative AI on learning remains limited, and much of it consists of short-term studies with small samples[4]. Nevertheless, existing mechanism-based analysis is sufficient to support a demanding judgement: cognitive symbiosis is not the natural product of human-machine interaction. It is an achievement premised on specific subject capacities and instructional design. Once these preconditions are absent, coupling falls from symbiosis into dependence. The three tensions discussed below are rooted precisely in this condition.

3. The Limits of Symbiosis: Three Inherent Tensions in Educational Digitalisation

The first tension appears at the level of individual cognition: augmentation and degradation often arise from the same source. Similar concerns were already visible in the era of search engines. Research identified the so-called "Google effect": when people believe that information can be retrieved at any time, they are more inclined to remember where information is located than to remember the information itself[28]. Generative AI extends what can be outsourced from memory to analysis, synthesis, and argumentation. A recent experiment points to "metacognitive laziness": although generative tools can help learners produce outputs more quickly, learners' monitoring of and reflection on their own learning processes may decline[29]. This is accompanied by over-reliance, in which machine outputs are accepted without critical scrutiny[30]. The risk is especially visible in writing. Writing is never merely the recording of pre-existing thoughts; it is itself a process of thinking. Through repeated work on wording and structure, people clarify vague intuitions and discover gaps in their reasoning. Carr[31] warned early on that technologies, while bringing convenience, also quietly reshape and may even make shallower the ways in which people think. When machines can write on students' behalf, the submitted text may look more polished, yet what is silently lost is the opportunity to think through writing. Kasneci et al.[32] are therefore concerned that large language models may encourage plagiarism and weaken the development of higher-order capacities such as critical thinking and independent expression. The

internal logic is not difficult to see: the very frictionlessness that makes AI convenient removes the cognitive friction on which capability grows.

The second tension appears at the level of social equity. Advocates often present AI as a powerful instrument for promoting fairness: it seems capable of making formerly expensive personalised tutoring available to all at almost zero marginal cost. Reality is more complex. At the level of access, the International Telecommunication Union[33] estimates that around 2.6 billion people worldwide remain offline, let alone able to use frontier AI services. More hidden are gaps at the levels of literacy and outcome. Whether AI can be transformed into an aid to symbiosis depends heavily on users´ self-regulation, critical thinking, and AI literacy, and these capacities themselves are unevenly distributed across groups with different socioeconomic backgrounds. The result may be a Matthew effect: those who already possess more cultural and cognitive capital are better able to use AI as a tool for self-enhancement, while disadvantaged learners are more likely, in the absence of guidance, to slide into dependence. Thus technology may not close gaps. It may instead, beneath the appearance of equal opportunity, create a deeper and less visible form of cognitive inequality[8-9]. In other words, the educational consequences of AI are often determined not by the availability of the technology but by the distribution of the capacity to use it.

The third tension appears at the level of knowledge. Personalisation is the most prominent promise of generative AI, yet beneath it lies the risk of homogenisation. The outputs of large language models are, in essence, probabilistic reproductions of statistical regularities in training corpora, and those corpora are not neutral in language, culture, or values. Bender, Gebru, McMillan-Major, and Shmitchell[34] describe such models as "stochastic parrots": they appear to generate meaning, but in fact stitch together existing patterns and may reproduce or even amplify biases and stereotypes embedded in them. When hundreds of millions of learners obtain knowledge, organise thought, and even craft expression through a small number of models, modes of expression and frameworks of thinking may converge, narrowing the cognitive diversity on which knowledge production depends. Still more troubling is the withdrawal of subjectivity. If students become accustomed to directly adopting the default answers offered by machines, their independent perspectives and judgements may be unconsciously replaced by homogenised outputs. How knowledge is organised and assessed in AI-enabled learning ecologies therefore needs to be reconsidered in its own right[35].

For learners in the non-English-speaking world, the risk of homogenisation is especially immediate. The training corpora of current mainstream large language models are dominated by English and Western knowledge traditions. Their implicit conceptual frameworks, value assumptions, and expressive paradigms may not fit the cultural contexts and educational aims of the Chinese-speaking world[34]. If learners organise their thoughts and acquire modes of expression through such models over long periods, local knowledge traditions and cultural subjectivity may be gradually marginalised. Cognitive homogenisation is therefore not only a cognitive problem but also a question of cultural diversity and educational sovereignty. Developing countries that are vigorously advancing educational digitalisation need to remain particularly alert to this risk. For precisely this reason, developing models and corpora suited to local contexts, and deliberately introducing multiple perspectives and critical comparison into teaching, is not merely a technical choice but a form of cultural and educational self-awareness.

These three tensions are located at the cognitive, social, and knowledge levels, but they share one source: when machine output replaces human cognitive labour and judgement, and this replacement fails to preserve human agency, augmentation turns into degradation, fairness turns into differentiation, and individuality turns into homogenisation. In sum, all three are transformations of cognitive symbiosis into cognitive dependence when its conditions are absent. The way forward is neither blind embrace nor rejection out of fear. It lies in systematically cultivating the conditions that can make human-machine interaction reliably tend towards

symbiosis. This is precisely the task that educational governance must undertake.

4. Towards Responsible Cognitive Symbiosis: Three Focal Points for Governance

Governance thinking must change accordingly: from controlling a "tool" to cultivating a "cognitive ecology". Its goal should not stop at expanding the scale and efficiency of application, but should protect and strengthen human agency and judgement in cognitive activity. In practice, governance clearly lags behind technology. A global survey by UNESCO[12] found that among hundreds of responding schools and universities, fewer than 10 per cent had developed institutional policies or formal guidance on generative AI. At the normative level, a new approach is consistent with the human-centred AI principles that have gained increasing consensus in recent years. The Beijing Consensus requires educational applications of AI to remain human-centred and to augment rather than replace human capacities[36]. Subsequent guidance further emphasises agency, inclusion, equity, and accountability[12][37]. Holmes and Tuomi[20] likewise call for pedagogy rather than technology to be placed at the centre of design. On this basis, responsible cognitive symbiosis can be established as a governance orientation and implemented through three mutually supporting dimensions: AI literacy, instructional design, and institutional safeguards.

The first focal point is to cultivate AI literacy beyond operational competence. Literacy is often narrowed to the ability to use tools, but in the age of generative AI it must include critical and metacognitive dimensions[38][39]. The critical dimension requires learners to understand the boundaries and inherent defects of models: they may confidently fabricate, a phenomenon commonly described as hallucination; they may embed biases; and their outputs should be examined rather than trusted uncritically. The metacognitive dimension requires learners to know clearly when they should think independently, when they may draw on machines, and how they can continue to retain their own judgement after doing so. Selwyn[40] argues that technology-oriented education should include three levels: learning with technology, learning about technology, and learning beyond technology. In this sense, education should not only teach people how to use AI, but also teach them to understand it and, when necessary, to think beyond it independently. Self-regulation and metacognitive capacity are the individual preconditions for cognitive symbiosis and should become core educational goals in the age of AI rather than optional additions.

The second focal point is to reconstruct instructional design and assessment. Since one condition of symbiosis is the preservation of desirable difficulty, instructional design should deliberately reserve spaces for thinking that cannot be outsourced and must be personally experienced by learners. Assessment must correspondingly shift from a focus on results to a focus on process. When a finished paper can be generated by a machine within seconds, judging learning solely by the final product has become almost ineffective. Assessment should move towards the evolution of ideas, the iteration of drafts, oral defence, and learners´ reflection on how they used AI. In the face of AI, many institutions´ first response has been to strengthen detection and block technology. Such containment is becoming increasingly ineffective technically and, more importantly, turns what should be a relationship of trust between teachers and students into an adversarial contest. It avoids the real question: if machines can already do the work, why should students still do it themselves? The way forward may be to shift from "preventing students from using AI" to "redefining learning that is worth assessing", and to include the capacity for appropriate AI use as an educational objective in its own right. The teacher´s role will also change accordingly: from transmitter of information, which is precisely what machines are good at, to designer of learning and guide of thinking, holding the educational direction within a human-in-the-loop form of collaboration.

The third focal point is to provide institutional safeguards. First, equity-oriented resource provision should make bridging the AI divide a priority. Through infrastructure investment, public services, and literacy training for disadvantaged groups, it should prevent technological

dividends from concentrating one-sidedly among already advantaged groups. Second, data and algorithm governance must protect data generated in learning processes, implement algorithmic accountability, and set prudent thresholds for applications aimed at minors. UNESCO[12], for example, recommends setting the minimum age for using generative AI tools in classrooms at 13. Governance should also vary by educational stage. In basic education, especially at younger ages, learners´ metacognition is still forming, and prudent restrictions are advisable to avoid crowding out the critical period for foundational capacity development. In higher education and professional training, where learners´ judgement is relatively more mature, the emphasis can shift from restriction to guidance. Finally, teacher professional development is essential, because even the best ideas must ultimately be implemented through teachers. This involves not only training in operational skills, but also helping teachers develop a critical understanding of AI and pedagogical judgement, so that they can use the strengths of technology while consciously protecting those educational processes in the classroom that cannot be outsourced. Teachers´ professional capacity is the key link connecting the governance measures described above. Only when teachers truly become the leaders of human-machine collaborative teaching rather than bystanders will the good intentions of institutional design avoid failing at the final point of implementation. Across all of these measures, governance must remain alert to a logic that prioritises technology and efficiency above all else. Educational outcomes can never be reduced to quantifiable indicators. More important than "doing things faster" is always "what kind of person is being formed".

It must also be emphasised that these three focal points are not isolated from one another. They need to be advanced collaboratively and require the joint participation of government, schools, technology companies, teachers, and families. The cultivation of individual literacy will fail without supporting changes in teaching and assessment. Instructional reform will be difficult to sustain without institutional resources. The absence of any one link may substantially weaken the overall effort. Particularly worthy of vigilance is a mismatch in governance tempo: technological iteration is measured in months, while educational institutions often adjust over years. This asymmetry of timescales means that governance cannot be limited to passive repair after problems have already emerged. It should establish more forward-looking and agile response mechanisms[20]. In this sense, responsible cognitive symbiosis is not a set of rules that can be issued once and for all. It is a dynamic process that must be continuously calibrated through ongoing interaction among technology, pedagogy, and institutions. The responsibility of educational researchers may lie precisely in continuing to ask what educational meanings technological applications carry, thereby providing rational reference points that are neither blindly optimistic nor reflexively resistant.

5. Conclusion

Some may argue that similar anxieties have appeared repeatedly in previous technological transformations: from Socrates´ concern as recorded in Plato´s Phaedrus that writing would corrupt memory, to the doubts surrounding the early spread of printing, calculators, and the Internet. Human minds, this reminder suggests, ultimately remained intact and even became richer. This historical reminder deserves to be taken seriously. It is an antidote to technological pessimism and warns us not to deny the potential of technology because of temporary discomfort. Yet there is a difference between generative AI and earlier technologies that cannot be ignored. Writing and printing externalised the storage of knowledge. Calculators externalised specific operations. Generative AI, however, can perform a series of higher-order activities at the core of human cognition: understanding, reasoning, and expression. When what is outsourced is thinking itself, the earlier optimistic analogy of tools extending capability is no longer self-evident. This does not mean that the conclusion must be pessimistic. It means that whether technology and mind can once again grow together in a desirable way will not happen automatically this time. It will depend on whether education responds appropriately and in time. We should be especially wary of a misleading comfort: the claim that, since students will ultimately coexist with AI, the earlier

and more extensively they depend on it, the better they will adapt to the future. This view confuses knowing how to use tools with being used by tools. Those who will stand securely in the age of AI are precisely those who have not allowed machines to think for them and who have thereby preserved and developed independent minds.

Generative AI is like a mirror, revealing those parts of education that cannot be replaced by algorithms. When retrieval, restatement, and even composition can be handed over to machines, education must ask with renewed urgency: what, in human growth, cannot be outsourced? The answer largely lies in capacities that machines cannot truly possess: judgement based on value weighing, meaning-making grounded in lived experience, and practical wisdom that integrates knowledge and virtue. There is also a more easily overlooked dimension: the relationality of education. Education is never merely the transmission of information. It is also an encounter between persons, grounded in understanding, trust, and care. Biesta[41] links the meaning of education to its "beautiful risk", its inability to be fully predicted and controlled, while Noddings[42] regards care as the ethical foundation of education. Generative AI can simulate dialogue with remarkable fluency, but it is difficult for it to be truly present or to offer understanding and recognition grounded in life. For this reason, cognitive symbiosis is not an inevitable gift of technological development, and cognitive dependence is not an unavoidable fate. The eventual direction depends on the choices we make in technological design, teaching practice, and institutional arrangement. The task of education in the age of AI is not so much to compete with machines as to use this disruption to reaffirm which capacities can only be cultivated by humans and passed from human beings to other human beings.

Authors' Biographies

Cao Hongyu was born in August 1979 in Langfang, Hebei Province, China. She is of Han ethnicity. She holds a Bachelor's degree and is affiliated with Langfang Zhenguigu Technology Co., Ltd. Her research interests include AI-empowered education and cognitive development.

Gan Weixiang was born in April 1995 in Nanchang, Jiangxi Province, China. He is of Han ethnicity. He holds a Ph.D. degree and is affiliated with Lincoln University College. His research interests include big data, higher education, and management science.

References:

[1] HU K. ChatGPT sets record for fastest-growing user base: analyst note[N]. Reuters, 2023-02-02.

[2] Digital Education Council. Digital Education Council global AI student survey 2024[R]. Digital Education Council, 2024.

[3] ZAWACKI-RICHTER O, MARÍN V I, BOND M, et al. Systematic review of research on artificial intelligence applications in higher education: where are the educators?[J]. International Journal of Educational Technology in Higher Education, 2019, 16(1): 39.

[4] CROMPTON H, BURKE D. Artificial intelligence in higher education: the state of the field[J]. International Journal of Educational Technology in Higher Education, 2023, 20(1): 22.

[5] ZHU Z, DAI L, ZHAO X. Opportunities, challenges, and development paths for educational applications of generative artificial intelligence[J]. China Educational Technology, 2023(7): 1-10.

[6] HOLMES W, BIALIK M, FADEL C. Artificial intelligence in education: promises and implications for teaching and learning[M]. Boston: Center for Curriculum Redesign, 2019.

[7] YU S. The future role of AI teachers[J]. Open Education Research, 2018, 24(1): 16-28.

[8] SELWYN N. Education and technology: key issues and debates[M]. 2nd ed. London:

Bloomsbury Academic, 2016.

[9] WILLIAMSON B, EYNON R. Historical threads, missing links, and future directions in AI in education[J]. Learning, Media and Technology, 2020, 45(3): 223-235.

[10] BEARMAN M, RYAN J, AJJAWI R. Discourses of artificial intelligence in higher education: a critical literature review[J]. Higher Education, 2023, 86(2): 369-385.

[11] GU X, DU H. Ethical risks and regulation of educational applications of artificial intelligence[J]. Journal of East China Normal University (Educational Sciences), 2022, 40(9): 1-12.

[12] UNESCO. Guidance for generative AI in education and research[R]. Paris: UNESCO, 2023.

[13] LUCKIN R, HOLMES W, GRIFFITHS M, et al. Intelligence unleashed: an argument for AI in education[M]. London: Pearson, 2016.

[14] VYGOTSKY L S. Mind in society: the development of higher psychological processes[M]. Cambridge: Harvard University Press, 1978.

[15] HUTCHINS E. Cognition in the wild[M]. Cambridge: MIT Press, 1995.

[16] PEA R D. Practices of distributed intelligence and designs for education[M]//SALOMON G. Distributed cognitions: psychological and educational considerations. Cambridge: Cambridge University Press, 1993: 47-87.

[17] CLARK A, CHALMERS D. The extended mind[J]. Analysis, 1998, 58(1): 7-19.

[18] SALOMON G, PERKINS D N, GLOBERSON T. Partners in cognition: extending human intelligence with intelligent technologies[J]. Educational Researcher, 1991, 20(3): 2-9.

[19] HEERSMINK R. A virtue epistemology of the Internet: search engines, intellectual virtues and education[J]. Social Epistemology, 2017, 31(1): 1-18.

[20] HOLMES W, TUOMI I. State of the art and practice in AI in education[J]. European Journal of Education, 2022, 57(4): 542-570.

[21] RISKO E F, GILBERT S J. Cognitive offloading[J]. Trends in Cognitive Sciences, 2016, 20(9): 676-688.

[22] BLOOM B S. The 2 sigma problem: the search for methods of group instruction as effective as one-to-one tutoring[J]. Educational Researcher, 1984, 13(6): 4-16.

[23] SWELLER J. Cognitive load during problem solving: effects on learning[J]. Cognitive Science, 1988, 12(2): 257-285.

[24] BJORK E L, BJORK R A. Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning[M]//GERNSBACHER M A, PEW R W, HOUGH L M, et al. Psychology and the real world. New York: Worth Publishers, 2011: 56-64.

[25] FLAVELL J H. Metacognition and cognitive monitoring: a new area of cognitive-developmental inquiry[J]. American Psychologist, 1979, 34(10): 906-911.

[26] ZIMMERMAN B J. Becoming a self-regulated learner: an overview[J]. Theory Into Practice, 2002, 41(2): 64-70.

[27] MOLLICK E. Co-intelligence: living and working with AI[M]. New York: Portfolio/Penguin, 2024.

[28] SPARROW B, LIU J, WEGNER D. Google effects on memory: cognitive consequences of having information at our fingertips[J]. Science, 2011, 333(6043): 776-778.

[29] FAN Y, TANG L, LE H, et al. Beware of metacognitive laziness: effects of generative artificial intelligence on learning motivation, processes, and performance[J]. British Journal of Educational Technology, 2025, 56(2): 489-530.

[30] BUÇINCA Z, MALAYA M B, GAJOS K Z. To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making[J]. Proceedings of the ACM on Human-Computer Interaction, 2021, 5(CSCW1): 1-21.

[31] CARR N. The shallows: what the Internet is doing to our brains[M]. New York: W. W. Norton, 2010.

[32] KASNECI E, SESSLER K, KÜCHEMANN S, et al. ChatGPT for good? On opportunities and challenges of large language models for education[J]. Learning and Individual Differences, 2023, 103: 102274.

[33] International Telecommunication Union. Measuring digital development: facts and figures 2023[R]. Geneva: ITU, 2023.

[34] BENDER E M, GEBRU T, MCMILLAN-MAJOR A, et al. On the dangers of stochastic parrots: can language models be too big?[C]//Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. New York: Association for Computing Machinery, 2021: 610-623.

[35] COPE B, KALANTZIS M, SEARSMITH D. Artificial intelligence for education: knowledge and its assessment in AI-enabled learning ecologies[J]. Educational Philosophy and Theory, 2021, 53(12): 1229-1245.

[36] UNESCO. Beijing consensus on artificial intelligence and education[R]. Paris: UNESCO, 2019.

[37] UNESCO. AI and education: guidance for policy-makers[R]. Paris: UNESCO, 2021.

[38] LONG D, MAGERKO B. What is AI literacy? Competencies and design considerations[C]// Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. New York: Association for Computing Machinery, 2020: 1-16.

[39] NG D T K, LEUNG J K L, CHU S K W, et al. Conceptualizing AI literacy: an exploratory review[J]. Computers and Education: Artificial Intelligence, 2021, 2: 100041.

[40] SELWYN N. Should robots replace teachers? AI and the future of education[M]. Cambridge: Polity Press, 2019.

[41] BIESTA G J J. The beautiful risk of education[M]. Boulder: Paradigm Publishers, 2013.

[42] NODDINGS N. Caring: a relational approach to ethics and moral education[M]. 2nd ed. Berkeley: University of California Press, 2013.