

Invariant Core, Changing Times: Rethinking Faculty Development for Teaching in the Generative AI Era

Shen Yong Ho

School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371, Singapore

© Higher Education Press 2026

Abstract With the accelerating tempo of technological change, universities' teaching reforms should be timely and principled rather than slow, ad hoc or reactive. This article anchors such reform in an invariant core: learning is underpinned by neuropsychological processes in which the cognitive and affective dimensions are inseparable. Using the Structure of Observed Learning Outcome taxonomy developed by Biggs as a conceptual lens, we argue that generative AI now produces artefacts exhibiting relational-level structure, and that the learner may bypass the deeper thinking that assignments are designed to elicit. We then discuss the evaluative judgments about cognitive offloading as both a risk and an opportunity in learning, guided by a principled test: Whether the process being bypassed is constitutive of the intended learning outcome. Three shifts in faculty development are proposed. First, development should deepen pedagogical knowledge and practice, placing greater emphasis on the mechanisms of human learning. Second, as AI-assisted coding lowers the barrier to creating educational tools, development should move from technology adoption towards pedagogical authoring of digital educational tools, anchored in the invariant core. Third, as our scientific understanding of learning develops, faculty should extend their scholarly identity as researchers to their teaching. They will also need to propose principled and measured adjustments to teaching approaches and curricula in step with the changing priorities of society. Faculty will have to be role models of ethical AI use and of ethical practices of their respective disciplines. The article closes with the reasons for learning and the institutional conditions universities must provide.

Keywords faculty development, generative AI, cognitive offloading, learning design, science of learning, SOLO taxonomy

1 Introduction

1.1 | Accelerating Tempo of Technological Change

In November 2022, ChatGPT brought conversational generative AI (GenAI) into mainstream public use (OpenAI, 2022). That arrival opened the very core of learning in educational institutions, the thinking process, to new possibilities while making cognitive offloading a pervasive concern (Lee et al., 2025; Risko & Gilbert, 2016). Just over three and a half years later, in June 2026, Google DeepMind released *From AGI to ASI*, a report mapping possible pathways from artificial general intelligence to artificial superintelligence (Genewein et al., 2026). The bar it sets is telling: Superintelligence is defined not as a system that surpasses an individual expert, but as one that is cognitively more capable than large organizations of humans.

The tempo of this change is striking. Education had several decades to adjust to the introduction of the internet, and only recently have governments moved to curb the harms of one of its offshoots: algorithmically driven social media for the young (Haidt, 2024).

Looking further back, the development from oral tradition to writing and then print took place over many centuries, each stage a slow restructuring of how knowledge is stored and transmitted (Eisenstein, 1979; Ong, 1982). Now the current seismic transformations take place in a matter of years or months, with skill sets needed for jobs after graduation rapidly evolving. How should educators respond, and what should faculty development look like in the future?

Given the above, this article proposes three major shifts in faculty development. First, inevitably, faculty will have to judge which forms of cognitive effort must be protected, what can be responsibly offloaded, and how the learning process can be evidenced. Thus, faculty development should give

Received July 14, 2026; revised September 9, 2026; accepted September 11, 2026

Shen Yong Ho (✉)

E-mail: hosing@ntu.edu.sg

greater attention to the underlying mechanisms of human learning as they do not change as quickly as technology and society. Second, as AI-assisted coding lowers the barrier to creating educational tools, faculty development should move beyond the adoption of educational technology towards pedagogical authorship: enabling faculty to design, adapt, and evaluate learning tools, guided by an understanding of the learning process and pedagogy, and informed by interpreted learning data. Third, as our scientific understanding of learning develops, faculty should extend their scholarly identity as researchers to their teaching, making principled curriculum adjustments as societal priorities shift, advising students on the responsible use of AI, and modelling both ethical AI practice and the ethical practices of their disciplines.

Yet across much of the education sector, change lags student adoption: Recent surveys consistently find student use of GenAI ahead of institutional guidance, tools, and support (Digital Education Council, 2024; Stephenson & Armstrong, 2026; UNESCO, 2025). The needed changes should be timely and principled rather than slow, ad hoc or reactive. We therefore begin with the orienting principle of the title, then examine what GenAI has changed about learning itself.

1.2 | An Orienting Principle: Anchor in the Core, Move with the Times

Here we propose a simple orienting principle, which we compress into eight Chinese characters of our own coinage, adapted from two classical Chinese sources: 不 离 于 中, 与 时 偕 行—anchor in the unchanging core (中, *zhong*), and move in step with the times. The first phrase adapts *Zhuangzi*’s “not departing from the ancestral source” (不 离 于 宗) (Ziporyn, 2020), substituting for 宗 (*zong*, the ancestral source) the character 中 (*zhong*), the unchanging, balanced core that lies at the root of all things (Johnston & Wang, 2012); the second echoes the *Yijing* commentaries’ injunction to “move together with the time” (Lynn, 1994). Ng (2017) identifies the same two-sided posture, “timely change, timeless constants,” as central to the Singapore

education system’s success. In that spirit, we propose an invariant core on which universities can anchor, and we examine the timely modifications that changing times demand.

An invariant core in education is that learning is underpinned by neuropsychological processes. Human learners attend, interpret, remember, reason, feel, value, persist, and regulate effort as they learn (Meltzoff et al., 2009; National Academies of Sciences, Engineering, and Medicine, 2018; National Research Council, 2000; Soderstrom & Bjork, 2015; Thomas et al., 2019). The cognitive and the affective dimensions are deeply intertwined (Immordino-Yang & Damasio, 2007; Pessoa, 2008) so teaching must engage both (Immordino-Yang, 2015; Pekrun, 2006). The changing priorities of society and employers (Anthropic, 2026; World Economic Forum, 2025), and the availability of technological tools for teaching and learning, can be disconcerting for faculty and universities. However, we should be reminded that the neuropsychological basis of learning does not change on technological or societal timescales. Thus, centred on the invariant core, we can evaluate whether available technologies enhance learning, or even develop such tools ourselves. Faculty will also have to judge how changing societal priorities will require curriculum and teaching approaches to be modified (Sawyer, 2006). See Figure 1.

Here, our discussion focuses on university education and faculty development. There are wider interactions between the university and the world; the interested reader can refer to Barnett (2018, 2024) on the ecological university and its eight interacting ecosystems. The three shifts proposed above develop this principle: The first belongs to the anchor, the second to the movement, and the third holds the two together.

1.3 | Crossing the Line: GenAI and the SOLO Taxonomy

The digitalization of information and its easy availability through the internet made information gathering easier. GenAI, however, automates the gathering, selection and deep synthesis of information on behalf of the user. This difference is best captured in Biggs’ Structure of Observed Learning Outcome (SOLO) taxonomy




Figure 1 Invariant core and changing times.

(Figure 2) (Biggs & Collis, 1982; Biggs & Tang, 2011). The internet, which allows efficient searches for information, made multistructural-level (surface learning) tasks easier, but for most well-designed relational-level (deep thinking) tasks, the learner still had to make the connections between information for themselves. Now, with well-crafted prompts, GenAI can produce artefacts exhibiting relational-level and even certain extended-abstract-level structure. Technology has crossed that line. The crucial point is that AI now produces the relational-level artefact while the learner may bypass the relational-level thinking that the assignment task was designed to elicit. What, then, are the cognitive impacts on learning? How should tasks be designed so that they still actively engage the mind in deep thinking within each disciplinary area to achieve the intended learning outcome?

One strategy is to make part of the information students work with non-digital, drawn from the social and physical world, making a complete synthesis harder to outsource to AI and more likely to require human judgement, embodied engagement, contextual interpretation, and interpersonal sense-making. This embodied learning (Immordino-Yang, 2015) and tacit learning (Polanyi, 1966) can bring the learning experience closer to real-world application. Two broad examples are a prototyping assignment (for Science, Technology, Engineering, and Mathematics (STEM)) or a task requiring human interaction (for the Arts, Humanities, and Social Sciences). In designing such activities, a deeper understanding of how people learn equips faculty to create the conditions, supports and challenges that help learners move from non-understanding to understanding (Bjork & Bjork, 2011; Meyer & Land, 2006). The SOLO taxonomy will return twice in what follows: as the scaffold for evaluative judgements on offloading, and as the frame for redesigning learning activities and assessment.

2 Invariant Core

This section establishes the anchor: the scientific basis

of the invariant core, the evaluative judgement it grounds, and the professional identity that carries such judgement.

2.1 | Learning as an Integrated Cognitive–Affective Process

Here we take up the first shift: grounding faculty development more deeply in the mechanisms of human learning. The discussion on the extent to which the learning sciences, the science of learning and neuroscience should be incorporated in faculty/teacher preparation is not new. The potential of these areas to provide insights into why certain pedagogical principles and teaching practices work well is undeniable (Dubinsky et al., 2013). However, the emphasis in faculty/teacher preparation typically falls on pedagogy and practice; these can be deepened by learning more about the underlying neuropsychological processes.

Learning sciences, science of learning, and neuroscience are often used loosely: A recent review found that researchers define the “science of learning” in strikingly varied ways and use it almost interchangeably with the “learning sciences” (Privitera et al., 2023). It therefore helps to set out what we mean by each. The science of learning includes the broad body of evidence on how people learn, drawing together cognitive psychology, neuroscience and education (Meltzoff et al., 2009). Neuroscience is narrower: It studies the brain itself and can inform teaching indirectly, helping to explain why a practice works rather than just showing that it does. It has too often been over-applied to the classroom, and claims of “brain-based” teaching frequently outrun the evidence (Howard-Jones, 2014; Thomas et al., 2019). The learning sciences are concerned with design: How we build the tasks, tools, and environments that put this understanding to use (Sawyer, 2006). Our invariant rests on all three. Its central claim, that thinking and emotion cannot be separated, comes from affective neuroscience


Pre-structural	Uni-structural	Multi-structural	Relational	Extended-abstract
				
No knowledge	Knows one relevant idea	Knows several relevant ideas separately	Integrates the ideas into a coherent whole	Generalizes integrated whole to new domains

Figure 2 Structure of Observed Learning Outcome (SOLO) taxonomy.

(Immordino-Yang & Damasio, 2007; Pessoa, 2008); turning that claim into real curricula and tasks is in the domain of learning sciences.

The affective dimension is easy to underrate, because so much of the discussion of AI and learning is framed in cognitive terms. Yet affect often shapes whether learners persist with the cognitive effort required for learning. Learners and teachers tend to judge how well they are learning by how fluent and easy the material feels; but this feeling is misleading. When students simply re-read material, their confidence in their own recall rises, even though their actual retention, tested after a delay, falls behind that of students who practised retrieving the material from memory instead (Roediger & Karpicke, 2006). The difficulty of retrieval is not a sign that learning is failing; it is often the sign that learning is working—a “desirable difficulty” (Bjork & Bjork, 2011). The trouble is that the difficulty feels like failure, and so learners, guided by that feeling, abandon the very effort that would have served them (Pekrun, 2006). Lodge (2025) captures this in an apt image: GenAI is the “e-bike of the mind,” carrying students efficiently to a finished assignment, while removing the exertion the journey requires. But in learning, unlike travel, the exertion is the key. The GenAI tool’s affective appeal, relief from mental strain, works directly against the cognitive process that the strain accompanies. The two strands of our invariant are here in tension: Upholding the cognitive requires resisting an affective impulse—to make things easier for students. Faculty who understand this can design for it—building in productive difficulty and helping students reinterpret that effort as necessary in real learning.

2.2 | Evaluative Judgement on Offloading Thinking

This returns us to GenAI’s cognitive impact on learning, and to the judgement at the heart of the first shift: what to protect and what to offload. Emerging evidence suggests that offloading thinking can reduce skill acquisition. The test is what learners can do unassisted. In a randomized field experiment with nearly a thousand high-school mathematics students, those given a standard ChatGPT-style assistant scored higher during practice, but on a subsequent unassisted exam they scored lower than students who had never used the assistant (Bastani et al., 2025). Notably, the harm largely disappeared when the assistant was redesigned to offer teacher-crafted hints rather than answers.

At larger scale, a 30-month panel study of 26,811 Chinese secondary students, reported in a working paper not yet peer-reviewed, found that AI adoption was associated with higher homework scores but lower subsequent unassisted exam performance

(Strömberg et al., 2026). The penalty on high-stakes entrance examinations emerged only after about two years, and was concentrated among the roughly 80% of users whose behaviour resembled homework outsourcing. While this is a caution to heed, the categorical view that offloading is undesirable for learning is not always correct (Gonsalves, 2026; Risko & Gilbert, 2016). In courses developing higher-order thinking skills, offloading tasks peripheral to the objective can free the learner for more complex challenges. Consider the same act of offloading in two different courses. In an introductory computer science course, having AI write the code bypasses the very skill the course is meant to develop, and so harms learning. On the other hand, in a prototyping course, using AI to code the hardware frees up time for students to explore more design possibilities, and so augments it. The more precise diagnostic, then, is not whether offloading occurs, but whether the effortful process being bypassed is itself constitutive of the intended learning outcome.

Through the lens of the SOLO taxonomy, in the introductory computer science course, knowing the constructs (e.g. conditionals, loops, functions) in isolation places the student at the multistructural level, but being able to integrate them into a coherent program or algorithm requires the student to operate at the relational level. Moving up to software design, the student will be at the deeper extended-abstract level. However, in a prototyping course, the programs, algorithms, hardware knowledge, and the context of the real-world problem become the entities at the multistructural level. It is integrating the programs and algorithms to control the hardware correctly to solve the real-world problem that makes the student operate at the relational level in the prototyping course, even with AI assistance in coding. Any ingenuity to generalise the context to a category of wider problems would be a demonstration of performing at the extended-abstract level. Gonsalves (2026) argues that Bloom’s taxonomy fails to address the cognitive demands of AI-assisted learning. Here the SOLO taxonomy remains a serviceable framework for analysing the level of thinking, even with AI use, because it classifies the structure of the observed outcome rather than a hierarchy of cognitive skills. Similar analysis can be extended to other disciplines, and for take-home assignments, the crossover to the relational and above levels will inevitably involve the physical and social world, echoing the strategy noted in the earlier section: a necessary but insufficient condition to support learners in engaging in deeper thinking, as AI becomes more powerful.

Since GenAI has crossed the line, the design of learning activities becomes more nuanced, and faculty must be equipped to make defensible evaluative judgements themselves in their course design (Bearman et al., 2024).

2.3 | Scholarly Teacher: Reflective Practice, Evaluative Judgement

Designing such learning activities well, however, is not a fixed competence. While we have identified “learning is underpinned by neuropsychological processes” as the invariant core, it is not a static core: Learning remains an active area of research. Faculty are already highly trained knowledge workers—that is the essence of the research role. What remains overdue is the extension of that professional identity to teaching: treating teaching not as content transmission but as a domain of scholarly practice in its own right (Boyer, 1990; Sawyer, 2006), subject to the same habits of evidence, reflection and revision that faculty apply in their disciplines. Such inquiry can culminate in findings made public and open to peer review through the scholarship of teaching and learning (Hutchings & Shulman, 1999).

In practice, this scholarly stance begins with the reflective teacher: one who treats each lesson as evidence, examining what worked and what did not (Schön, 1983) and viewing that practice through the eyes of students, colleagues and research literature (Brookfield, 2017). Here, technology itself can serve that reflection: AI-enhanced software such as Engageli and ClassIn captures the lesson on video with fine-grained interaction data, and is beginning to make it possible for AI to surface patterns in teaching practice (Blikstein & Worsley, 2016).

The same scholarly judgement must extend beyond the technologies of classroom teaching to the wider technological environment in which students learn. As the case of social media has shown, the way a technology is used shapes its impact on the young—and that impact unfolds largely outside formal instruction (Haidt, 2024). GenAI now permeates this environment: Students encounter it in their studying, assignments and daily lives. Faculty should therefore develop the evaluative judgement to decide how, and whether, a technology should be used in their courses, grounded in their understanding of the neuropsychological processes of learning.

3 Changing Technologies and Changing Societal Priorities

This section turns to the movement: Authoring learning tools anchored in the core, adjusting courses and curricula as societal priorities change, and tracing a short history of teaching programming as an extended illustration.

3.1 | Faculty Members as Technological Learning Designers

The second shift is pedagogical authorship. GenAI

differs from earlier educational technologies in one important respect. Its simplest use is as a chatbot for students, with responses grounded in faculty-selected content and behaviour (such as requiring Socratic dialogue) pre-determined through faculty-designed meta-prompts (Bastani et al., 2025; Mollick & Mollick, 2024). Beyond this, it also allows a non-programmer faculty member to prompt-code (more commonly known as *vibe coding* (Collins Dictionary, 2025)) computer programs for educational purposes. Agentic development environments, such as Google’s Antigravity at the time of writing (Google Antigravity Team, 2025), allow a user to describe an intended application in natural language and have AI agents plan, write and test the working software. In our own exploratory trial, a simple three-level physics game on projectile motion was prompted and iteratively enhanced (in about twenty prompts) within two hours, ready for classroom piloting. More important than the speed is that the process allows the faculty member to infuse pedagogical decisions into the design, scaffolding the learning and providing encouragement as needed (Koedinger & Aleven, 2007; Wood et al., 1976). Design judgements that once resided with the developers of specialized tutoring systems can now rest with the teacher.

This is the ideal versatility, where the programmer is also the user. A caveat is necessary. However, agentic tools may request extensive file-system access and the ability to execute commands, so isolation, such as a separate machine or a sandboxed virtual environment, should be considered. For online deployment, prompting the agents to incorporate cyber-security protections is a sensible start but not a sufficient one. Tools that collect student data should also pass through the institution’s usual IT-security and data-protection review.

As argued in the section on crossing the line, the artefact alone no longer evidences the thinking. With easy access to AI, the weighting of assessment must therefore shift from end-products (e.g., essays, posters, and presentations) to the process of developing them. Evaluating processes requires extended observation and data collection, something severely constrained by limited manpower. Here, prompt-coding gives faculty the versatility to develop platforms through which learning processes can be tracked and evaluated. For example, a system that captures in-class supervised writing, online conversations, iterative feedback and out-of-class writing can help shape more productive learning behaviour and, structured using the SOLO taxonomy, can help faculty align intended learning outcomes with learning activities and assessment (Biggs & Tang, 2011).

As faculty discover the ease of designing software for learning in natural language, and of collecting and analysing student learning data (Blikstein &

Worsley, 2016), it becomes even more important that these designs are grounded in the science of learning and sound pedagogy—our invariant core. Digitalization makes certain pedagogical practices that personalise learning easier to scale. For example, students can role-play with an AI chatbot in specified scenarios (Mollick & Mollick, 2024). Learning by teaching can be deployed through teachable AI agents (Chase et al., 2009; Roscoe & Chi, 2007). Some of these agents are engineered to express emotion and motivation, to engage the learner affectively (Borjigin et al., 2011; Zhao et al., 2012). Two trade-offs should be noted. First, in the role-playing and learning-by-teaching examples, the students are not communicating with another human. Designers may accept this trade-off for a system that is ubiquitous and gives instant, automated feedback (Immordino-Yang & Damasio, 2007). Second, a subtler issue is that certain forms of personalization, in which the system is designed to smooth away difficulty for the individual student, risk making learners less agile and less adaptable: The conditions that ease immediate performance are often not the conditions that build durable, transferable learning (Bjork & Bjork, 2011; Soderstrom & Bjork, 2015). These, again, are nuances the faculty learning designer must weigh (Bearman et al., 2024).

3.2 | Faculty as Course and Curriculum Designer

The movement extends beyond individual tools to the course and the curriculum. AI's breadth of knowledge and this versatility in developing learning tools now make possible more tailored ways of teaching and learning. Yet not all learning with AI takes place within settings the teacher directly designs. The faculty role thus extends beyond design to advising students on the use of AI, and to modelling not only ethical and epistemically honest AI practice but the ethical practices of the discipline itself: its research integrity, standards of evidence, and professional norms, into which students are being inducted (Cruess et al., 2008). Faculty are thus well placed to close the gap between student development under a traditional curriculum and the changing priorities of society and employers. This raises another tension, between what learning should entail and what the changing priorities require (Ashwin, 2020). The university is also an institution of knowledge discovery and creation. Responsiveness to changing priorities must be balanced against that identity, lest the institution of higher learning drift into an institute of vocational training. This is a worthy mission, but a different one.

Typical curricula are already packed: in addition to their subject core, students must equip themselves with new skills for new technologies and be

versed in broader interdisciplinary subjects to deal with real-world problems (OECD, 2020). Often this involves removing or condensing traditional content. To illustrate, the reader may attempt a simple exercise: Should schools still teach children to read analogue clocks, now that digital displays are everywhere? Pause on your own answer before reading on. Beyond telling time, reading the analogue clock is probably a child's first encounter with simple ambiguity and its resolution: determining the hour when the hour hand lies between two numbers. Other embedded concepts include fractions, multiples of five, sub-divisions, relating time to day and night, angles, mapping a continuously recurring quantity onto a circle, clockwise and anticlockwise. The list goes on. What appears to be an obsolete surface skill turns out to be a carrier of constitutive learning—the same diagnostic we applied to offloading, now at the scale of curriculum content. Curriculum reduction is thus akin to a game of Jenga: Some blocks are load-bearing in ways not visible from the outside. Yet reduction is a necessity as new job skills emerge, and GenAI has not only crossed the line but also raised the bar—As the next section, tracing how the teaching of computer programming has evolved, will illustrate. Readers can extend the exercise to imagine how their own disciplines and skills will be affected, at an accelerated pace, in the era of AI.

3.3 | How Computer Programming Curriculum Has Evolved Over Time

The rapid development of digital technology over the past sixty years has had arguably its starkest curricular impact in computer science itself. It has also reached into every university degree. Consider how the teaching of machine- and assembly-level programming, where human code communicates directly with the hardware, has evolved. As compilers matured and production programming moved to high-level languages, hand-coded assembly receded from professional practice. The displacement was already visible within a year of Fortran's release (Backus, 1978). The curriculum moved in step. Assembly programming went from a required course in the degree core (Atchison et al., 1968; Austing et al., 1979) to a section within computer architecture and systems courses today, taught for systems literacy rather than production (Bryant & O'Hallaron, 2001; Joint Task Force on Computing Curricula et al., 2013). What was once a required skill is now an awareness. The programming languages taught have likewise evolved, from Fortran and BASIC through C and Java to Python. The language debate itself receded as curricula converged on concepts common to all languages (Becker & Quille, 2019).

In the 1960s and 1970s, outside science and engineering, programming was not a required skill. The

interesting exception was Dartmouth, where computing was framed as essential to liberal education and required of most freshmen through their mathematics courses (Kemeny & Kurtz, 1968). Society has since evolved. Work and daily life have become digital across every sector, and this societal change is reflected in university-wide curricula. For example, in Singapore, universities have in recent years instituted computational, digital or data literacy components for their undergraduates, including those in the humanities, arts and social sciences, with each institution calibrating the requirement to its own mission. At Nanyang Technological University, the author's institution, digital-literacy core modules were introduced for all undergraduates from 2018 under the university's Smart Campus initiative (Nanyang Technological University, 2018). From the academic year 2021/22 intake, the interdisciplinary collaborative core added, among its common courses, a foundational course on the impact of technology, including AI, on humanity. The course is taken by every undergraduate in their first two years, and situates computational thinking alongside ethics, communication and inquiry, and global challenges (Nanyang Technological University, n.d.). This is an institutional expression of a connected curriculum that links disciplines to one another and outward to the world (Fung, 2017). Notably, what has been made universal is computational thinking rather than any particular language or tool: Programming itself remains discipline-specific, required where it is constitutive of the field of study. Faculty will need the skills to weave such courses meaningfully into the context of each degree programme, grounded in the science of learning and informed by the priorities and demands of society and employers (Anthropic, 2026; World Economic Forum, 2025). Such reports provide system-level signals. It is equally important that some faculty maintain continual communication with alumni, employers, and the wider society, to appreciate at first hand the changes needed in the training of their students. The bar continues to rise. As coding in natural language becomes possible with AI, the demand shifts upward to software design: to specifying what should be built and judging what is produced (Denny et al., 2024).

4 Conclusions: Scholarly Teacher

The three shifts converge on one figure: the scholarly teacher. The literature has shown that academic development involves the formation of teacher identity (van Lankveld et al., 2017), participation in communities, and significant conversations about teaching (Cox, 2004; Roxå & Mårtensson, 2009), reflective and scholarly inquiry into practice (Hutchings & Shulman, 1999; Schön, 1983), attention to disciplinary ways of knowing

and signature pedagogies (Shulman, 2005), and engagement with institutional cultures that enable or constrain teaching improvement (Trowler & Cooper, 2002). These traditions remain valuable, but GenAI changes the focus of development.

Our understanding of the invariant core evolves with research, and the changing layer now moves on a frontier already showing early signs of reshaping entry-level employment (Brynjolfsson et al., 2026). Development must therefore be continual. It must also equip faculty to read societal signals as design inputs, both from system-level reports and through the first-hand channels of communication noted in the programming illustration. Faculty evaluation must keep pace as well: Polished teaching artefacts (such as syllabi, slides, and rubrics) are subject to the same collapse as student artefacts. One symmetric caution should be kept in view: The offloading diagnostic applies to faculty and to development programmes themselves. A programme that uses AI to generate its own materials faces exactly the question it trains faculty to ask; the invariant applies to everyone who learns.

Yet none of this growth can rest on exhortation alone; it requires institutional conditions. A robust definition of teaching excellence gives faculty a clear standard to grow toward. An evaluation system that recognizes quality teaching and motivates meaningful, timely change ensures that growth is seen and rewarded. To do so, it must probe design reasoning directly: Why a task was engineered as it was, what was given and what withheld, and how the design serves the invariant core. A professional development programme that itself evolves keeps the anchoring and the moving within reach of every faculty member. A university that asks its faculty to anchor in the invariant core and move with the times must also ask who pays for and protects that movement, for the continued flourishing of the university is an institutional achievement, not an individual one. Under these conditions, faculty development anchors faculty in the invariant core and equips them to make timely, local evaluative judgements as technologies, disciplines and priorities change. The outcome is the scholarly teacher: reflective in practice, designing with evidence, and reading the signals of a changing society.

We began with an accelerating tempo. It is natural to close by asking what such a trajectory leaves for human learning. This returns us to what Biesta (2009) calls the question of purpose in education. The record of earlier "superhuman" moments suggests an answer. AlphaGo, an earlier landmark in modern deep-learning systems, defeated the European Go champion (Silver et al., 2016). Yet people continue to learn, study and teach the game. An analysis of 5.8 million professional moves across seven decades found that the quality of human play improved after the advent of superhuman

AI (Shin et al., 2023). Novel moves became more frequent, and more strongly associated with improved decision quality. Machine assistance has entered even pure mathematics. AI first guided mathematicians toward new conjectures and theorems (Davies et al., 2021), then made new discoveries for established open problems (Romera-Paredes et al., 2024). Most recently, in preprint results not yet peer-reviewed, AI has autonomously resolved several long-open Erdős problems, two of them open for over half a century, with machine-verified formal proofs (Tsoukalas et al., 2026). Yet students still recreate classical proofs by hand, because the eventual proof on the page was never the point; the skills developed in constructing it are. We have seen the analogue clock: an apparently obsolete artefact that turns out to be a dense carrier of constitutive learning. In each case, when technology takes over the production of the artefact, the reason for the human activity does not vanish. It migrates from production of artefacts to cultivation of the mind.

Acknowledgments No funding was received to assist with the preparation of this manuscript.

Declaration on the Use of AI During the preparation of this article, the author used Claude (Anthropic; Claude Fable 5 model) and ChatGPT (OpenAI; ChatGPT 5.6 Pro) to verify literature and citations against publisher and primary sources, critically evaluate and refine the author's drafted text, draft selected passages under the author's direction, and assist with document preparation. The author reviewed and edited all output and takes full responsibility for the content of the published article.

Author Contributions Shen Yong Ho conceived the article, developed the argument, and wrote and revised the manuscript.

Conflicts of Interest The author declares no conflict of interest.

Data Availability Statements Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

References

- Anthropic. (2026, June 26). *The Anthropic economic index*. Available from Anthropic website.
- Ashwin, P. (2020). *Transforming university education: A manifesto*. Bloomsbury Academic.
- Atchison, W. F., Conte, S. D., Hamblen, J. W., Hull, T. E., Keenan, T. A., Kehl, W. B., McCluskey, E. J., Navarro, S. O., Rheinboldt, W. C., ... Young, D. M., Jr. (1968). Curriculum 68: Recommendations for academic programs in computer science. *Communications of the ACM*, 11(3), 151–197.
- Austing, R. H., Barnes, B. H., Bonnette, D. T., Engel, G. L., & Stokes, G. (1979). Curriculum '78: Recommendations for the undergraduate program in computer science—A report of the ACM curriculum committee on computer science. *Communications of the ACM*, 22(3), 147–166.
- Backus, J. (1978). The history of FORTRAN I, II, and III. *ACM SIGPLAN Notices*, 13(8), 165–180.
- Barnett, R. (2018). *The ecological university: A feasible utopia*. Routledge.
- Barnett, R. (2024). *Realizing the ecological university: Eight ecosystems, their antagonisms and a manifesto*. Bloomsbury.
- Bastani, H., Bastani, O., Sungu, A., Ge, H. S., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences of the United States of America*, 122(26), e2422633122.
- Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment & Evaluation in Higher Education*, 49(6), 893–905.
- Becker, B. A., & Quille, K. (2019). 50 years of CS1 at SIGCSE: A review of the evolution of introductory programming education research. In: *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*. Minneapolis: ACM, 338–344.
- Biesta, G. (2009). Good education in an age of measurement: On the need to reconnect with the question of purpose in education. *Educational Assessment, Evaluation and Accountability*, 21(1), 33–46.
- Biggs, J. B., & Collis, K. F. (1982). *Evaluating the quality of learning: The SOLO taxonomy*. Academic Press.
- Biggs, J., & Tang, C. (2011). *Teaching for quality learning at university: What the student does*. 4th ed. Open University Press.
- Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In: Gernsbacher, M. A., Pew, R. W., Hough, L. M., & Pomerantz, J. R., eds. *Psychology and the real world: Essays illustrating fundamental contributions to society*. Worth, 56–64.
- Blikstein, P., & Worsley, M. (2016). Multimodal learning analytics and education data mining: Using computational technologies to measure complex learning tasks. *Journal of Learning Analytics*, 3(2), 220–238.
- Borjigin, A., Shen, Z. Q., & Miao, C. Y. (2011). Affective teachable agent in VLE: A goal oriented approach. In: *Proceedings of the 2011 IEEE 11th International Conference on Advanced Learning Technologies*. Athens: IEEE, 110–114.
- Boyer, E. L. (1990). *Scholarship reconsidered: Priorities of the professoriate*. Carnegie Foundation for the Advancement of Teaching.
- Brookfield, S. D. (2017). *Becoming a critically reflective teacher*. 2nd ed. Jossey-Bass.
- Bryant, R. E., & O'Hallaron, D. R. (2001). Introducing computer systems from a programmer's perspective. In: *Proceedings of the 32nd SIGCSE Technical Symposium on Computer Science Education*. Charlotte: ACM, 90–94.

- Brynjolfsson, E., Chandar, B., & Chen, R. (2026, August 12). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Available from Stanford Digital Economy Lab website.
- Chase, C. C., Chin, D. B., Oppizzo, M. A., & Schwartz, D. L. (2009). Teachable agents and the protégé effect: Increasing the effort towards learning. *Journal of Science Education and Technology*, 18(4), 334–352.
- Collins Dictionary. (2025, November 6). *Collins' Word of the Year 2025: AI meets authenticity as society shifts*. Available from Collins Dictionary Language Blog website.
- Cox, M. D. (2004). Introduction to faculty learning communities. *New Directions for Teaching and Learning*, 2004(97), 5–23.
- Cruess, S. R., Cruess, R. L., & Steinert, Y. (2008). Role modelling—Making the most of a powerful teaching strategy. *BMJ*, 336(7646), 718–721.
- Davies, A., Veličković, P., Buesing, L., Blackwell, S., Zheng, D., Tomašev, N., Tanburn, R., Battaglia, P., Blundell, C., ... Kohli, P. (2021). Advancing mathematics by guiding human intuition with AI. *Nature*, 600(7887), 70–74.
- Denny, P., Prather, J., Becker, B. A., Finnie-Ansley, J., Hellas, A., Leinonen, J., Luxton-Reilly, A., Reeves, B. N., Santos, E. A., & Sarsa, S. (2024). Computing education in the era of generative AI. *Communications of the ACM*, 67(2), 56–67.
- Digital Education Council. (2024, August 2). *Digital Education Council global AI student survey 2024*. Available from Digital Education Council website.
- Dubinsky, J. M., Roehrig, G., & Varma, S. (2013). Infusing neuroscience into teacher professional development. *Educational Researcher*, 42(6), 317–329.
- Eisenstein, E. L. (1979). *The printing press as an agent of change: Communications and cultural transformations in early-modern Europe*. Cambridge University Press.
- Fung, D. (2017). *A connected curriculum for higher education*. UCL Press.
- Genewein, T., Franklin, M., Lerchner, A., Orseau, L., Albanie, S., Bales, A., Wyeth, C., Chan, S., Gabriel, I., ... Legg, S. (2026). From AGI to ASI. *arXiv Preprint*, arXiv:2606.12683.
- Gonsalves, C. (2026). Generative AI's impact on critical thinking: Revisiting Bloom's taxonomy. *Journal of Marketing Education*, 48(1), 4–19.
- Google Antigravity Team. (2025, November 20). *Build with Google Antigravity, our new agentic development platform*. Available from Google Developers Blog website.
- Haidt, J. (2024). *The anxious generation: How the great rewiring of childhood is causing an epidemic of mental illness*. Penguin Press.
- Howard-Jones, P. A. (2014). Neuroscience and education: Myths and messages. *Nature Reviews Neuroscience*, 15(12), 817–824.
- Hutchings, P., & Shulman, L. S. (1999). The scholarship of teaching: New elaborations, new developments. *Change: The Magazine of Higher Learning*, 31(5), 10–15.
- Immordino-Yang, M. H. (2015). *Emotions, learning, and the brain: Exploring the educational implications of affective neuroscience*. W. W. Norton.
- Immordino-Yang, M. H., & Damasio, A. (2007). We feel, therefore we learn: The relevance of affective and social neuroscience to education. *Mind, Brain, and Education*, 1(1), 3–10.
- Johnston, I., & Wang, P. (2012). *Daxue and Zhongyong: Bilingual edition*. The Chinese University Press.
- Joint Task Force on Computing Curricula, Association for Computing Machinery (ACM), & IEEE Computer Society. (2013). *Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science*. ACM Press.
- Kemeny, J. G., & Kurtz, T. E. (1968). Dartmouth time-sharing. *Science*, 162(3850), 223–228.
- Koedinger, K. R., & Aleven, V. (2007). Exploring the assistance dilemma in experiments with cognitive tutors. *Educational Psychology Review*, 19(3), 239–264.
- 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*. Yokohama: ACM, 1–22.
- Lodge, J. (2025, February 14). *Generative AI is the e-bike of the mind*. Available from Future Campus website.
- Lynn, R. J. (1994). *The classic of changes: A new translation of the I Ching as interpreted by Wang Bi*. Columbia University Press.
- Meltzoff, A. N., Kuhl, P. K., Movellan, J., & Sejnowski, T. J. (2009). Foundations for a new science of learning. *Science*, 325(5938), 284–288.
- Meyer, J. H. F., & Land, R. (2006). *Overcoming barriers to student understanding: Threshold concepts and troublesome knowledge*. Routledge.
- Mollick, E., & Mollick, L. (2024). Instructors as innovators: A future-focused approach to new AI learning opportunities, with prompts. *arXiv Preprint*, arXiv:2407.05181.
- Nanyang Technological University. (2018, January 10). *NTU president Subra Suresh unveils plans to make NTU a smart campus*. Available from Nanyang Technological University website.
- Nanyang Technological University. (n.d.). *Interdisciplinary collaborative core (ICC)*. Available from Nanyang Technological University website.
- National Academies of Sciences, Engineering, and Medicine. (2018). *How people learn II: Learners, contexts, and cultures*. National Academies Press.
- National Research Council. (2000). *How people learn: Brain, mind, experience, and school: Expanded edition*. National Academies Press.
- Ng, P. T. (2017). *Learning from Singapore: The power of paradoxes*. Routledge.
- OECD. (2020). *Curriculum overload: A way forward*. OECD Publishing.
- Ong, W. J. (1982). *Orality and literacy: The technologizing of the word*. Methuen.
- OpenAI. (2022, November 30). *Introducing ChatGPT*. Available from OpenAI website.

- Pekrun, R. (2006). The control-value theory of achievement emotions: Assumptions, corollaries, and implications for educational research and practice. *Educational Psychology Review*, 18(4), 315–341.
- Pessoa, L. (2008). On the relationship between emotion and cognition. *Nature Reviews Neuroscience*, 9(2), 148–158.
- Polanyi, M. (1966). *The tacit dimension*. Doubleday.
- Privitera, A. J., Ng, S. H. S., & Chen, S. H. A. (2023). Defining the science of learning: A scoping review. *Trends in Neuroscience and Education*, 32, 100206.
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688.
- Roediger, H. L., III, & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255.
- Romera-Paredes, B., Barekatin, M., Novikov, A., Balog, M., Kumar, M. P., Dupont, E., Ruiz, F. J. R., Ellenberg, J. S., Wang, P. M., ... Fawzi, A. (2024). Mathematical discoveries from program search with large language models. *Nature*, 625(7995), 468–475.
- Roscoe, R. D., & Chi, M. T. H. (2007). Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions. *Review of Educational Research*, 77(4), 534–574.
- Roxå, T., & Mårtensson, K. (2009). Significant conversations and significant networks-exploring the backstage of the teaching arena. *Studies in Higher Education*, 34(5), 547–559.
- Sawyer, R. K. (2006). *The Cambridge handbook of the learning sciences*. Cambridge University Press.
- Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.
- Shin, M., Kim, J., van Opheusden, B., & Griffiths, T. L. (2023). Superhuman artificial intelligence can improve human decision-making by increasing novelty. *Proceedings of the National Academy of Sciences of the United States of America*, 120(12), e2214840120.
- Shulman, L. S. (2005). Signature pedagogies in the professions. *Daedalus*, 134(3), 52–59.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., ... Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
- Soderstrom, N. C., & Bjork, R. A. (2015). Learning versus performance: An integrative review. *Perspectives on Psychological Science*, 10(2), 176–199.
- Stephenson, R., & Armstrong, C. (2026). *Student generative AI survey 2026*. HEPI Report 199. Higher Education Policy Institute.
- Strömberg, D., Lei, V., & Wu, Y. H. (2026). *DP21577 the generative AI learning penalty: Evidence from Chinese secondary education*. CEPR Discussion Paper No. 21577. CEPR Press.
- Thomas, M. S. C., Ansari, D., & Knowland, V. C. P. (2019). Annual research review: Educational neuroscience: Progress and prospects. *Journal of Child Psychology and Psychiatry*, 60(4), 477–492.
- Trowler, P., & Cooper, A. (2002). Teaching and learning regimes: Implicit theories and recurrent practices in the enhancement of teaching and learning through educational development programmes. *Higher Education Research & Development*, 21(3), 221–240.
- Tsoukalas, G., Kovsharov, A., Shirobokov, S., Surina, A., Firsching, M., Bérczi, G., Ruiz, F. J. R., Suggala, A., Wagner, A. Z., ... Chaudhuri, S. (2026). Advancing mathematics research with AI-driven formal proof search. *arXiv Preprint*, arXiv:2605.22763.
- UNESCO. (2025, September 2). *UNESCO survey: Two-thirds of higher education institutions have or are developing guidance on AI use*. Available from UNESCO website.
- van Lankveld, T., Schoonenboom, J., Volman, M., Croiset, G., & Beishuizen, J. (2017). Developing a teacher identity in the university context: A systematic review of the literature. *Higher Education Research & Development*, 36(2), 325–342.
- 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.
- World Economic Forum. (2025). *The future of jobs report 2025*. Geneva: World Economic Forum.
- Zhao, G. P., Borjigin, A., & Shen, Z. Q. (2012). Learning-by-teaching: Designing teachable agents with intrinsic motivation. *Educational Technology & Society*, 15(4), 62–74.
- Ziporyn, B. (2020). *Zhuangzi: The complete writings*. Hackett Publishing.