Skills Never Formed: AI Adoption and Career Formation Risk in IT Higher Education

Dinko Horvat
Zagreb University of Applied Sciences
Vrbik 8, 10000 Zagreb, Croatia
dinko.horvat@tvz.hr

Vjeran Bušelić
University of Zagreb
Faculty of Organization and Informatics
Pavlinska 2, 42000 Varaždin, Croatia
vbuselic@foi.hr

***Abstract*[1].** *Generative AI arrived beneath educational institutions rather than through them, ambient in everyday tools. Computing students face double exposure: the same technology restructures both their learning environment and the profession they are entering. This conceptual synthesis integrates cognitive-science learning theory with three convergent randomized trials to trace one mechanism — how AI use at the formation stage, without metacognitive discipline, produces never-skilling: the failure to acquire professional capabilities because AI performs the very work through which they would form. From it the paper derives hold–rethink–continue, a teachable metacognitive posture. A companion paper develops the educator response.*

Keywords. generative AI, IT education, never-skilling, double exposure, hold–rethink–continue, schema formation, self-regulated learning, verification economy

1 Introduction

IT students enter higher education at the start of a dynamic and demanding career — in software development, data analysis, infrastructure, cybersecurity, and related fields. AI is already part of that world. It just appeared as an ambient layer already in the editor and the browser, before anyone decided how it should be used in business, society, and education. And it is changing two things at once: how they learn, and how it is rapidly changing the industry they are educated to join — a double exposure (adapting O'Brien and Leichenko's [2000] framework, originally two distinct global forces, to a single technology on two fronts) that no other field of study faces in the same form.

This paper is not against AI. Students already use it, and they need to — working with AI is now part of the profession, and learning to do it well is part of their education. The real question is how they adapt: whether AI becomes a tool they learn with, or one that does the learning for them.

The risk is specific. During the formation stage — the early period when professional capability is built — AI can complete the very tasks through which that capability would form, so a student can produce credible work without developing the independent skill behind it. The credential — the diploma — forms; the capability it implies stays in question. This paper calls that outcome never-skilling — a term now emerging across professional-education research (e.g., Ke et al., 2026, in medical education) for the failure Macnamara et al. (2024) describe as AI-induced skill development hindrance. It is not a sharper form of deskilling but a totally different risk.

How never-skilling differs from deskilling — in when it strikes, its mechanism, whether it can be detected, and what remedy it admits — is set out in Section 2, where the paper's scope is defined (Table 1.).

The concern is neither speculative nor settled. A cluster of 2025–2026 studies has begun to measure how AI use during learning affects the competence that forms beneath the visible output (Shen & Tamkin, 2026; Bastani et al., 2025; Fan et al., 2025). What that literature has barely examined is the population most exposed to the risk — computing students at the formation stage, for whom AI restructures the learning environment and the profession at once. This paper addresses that gap from the student's standpoint; Section 2 sets out its approach and Section 4 the evidence it rests on.

The remainder of this section sets out why this risk is acute now: the gap between computing curricula and industry (Section 1.1), the labour-market pressure that transmits it (Section 1.2), and why generative AI differs from earlier tools (Section 1.3). Section 2 sets out the paper's approach and scope; Section 3 traces the cognitive risk cascade that produces never-skilling; Section 4 presents the empirical evidence; Section 5 proposes student-facing practices; Section 6 discusses

[1] This paper is one of a pair developed in parallel for this conference. It addresses formation-stage AI adoption from the IT student's perspective. A companion paper Schema Before Output: Metacognitive Schema-First Scaffolding in Computing Higher Education in the GenAI Era addresses the institutional and pedagogical response. Each paper stands independently.

the findings and their limitations; and Section 7 concludes.

1.1 The Curricula–Industry Gap

For years, computing degrees were built for a job market that is now disappearing — one where new graduates spent their first two to four years writing functions, debugging failures, and building components from specifications (CC2020 Task Force, 2020). It was during their junior phase, not in the classroom, that developers built their core skills. These are the patterns and habits experienced developers use almost automatically. Now, generative AI speeds up that stage, since boilerplate and scaffolding can be created in seconds. As a result, the process that once helped new developers build good judgment is getting shorter, if it survives at all. The World Economic Forum (2025) records the pace: information and technology services report the highest expected AI transformation of any sector (99%), alongside declining demand for routine programming. Curricula cannot keep up: revision cycles run three to eight years against an adoption curve that accelerated between 2022 and 2025. The result is a structural mismatch. Graduates are expected to step straight into orchestration work — problem framing, system design, quality assurance of AI output — that presupposes the very judgment the compressed junior years used to build (Gartner, 2025).

1.2 Labour-Market Pressure

That pressure is sharpest in software development, where adoption has moved fastest: 84% of surveyed developers use or plan to use AI tools, with 51% of professional developers now using them daily (Stack Overflow, 2025), and Gartner (2025) projects 90% enterprise adoption by 2028. It reaches students through job postings that demand AI proficiency and cohort dynamics that penalise whoever abstains — and it arrives at the formation stage, before students have the schemas that would make AI assistance adaptive rather than substitutive (Prather et al., 2023). The result is an adoption dilemma: adopt early and gain a hiring edge at the cost of formation, or resist and protect formation at a real competitive cost. Neither response is irrational, which is precisely why the tension is structural — it calls for an institutional response, not individual resolve.

1.3 Why AI Is Different from Prior Productivity Tools

Offloading cognitive work to tools is neither new nor harmful in itself — calculators offload arithmetic, compilers offload syntax. What matters is **what** gets offloaded, and **when**. Generative AI offloads a different class of task: problem decomposition, pattern recognition, error diagnosis, solution generation. Every earlier tool changed what a developer worked **on**; this one offers to absorb what a developer thinks **with**. In professional work with established schemas that is a legitimate efficiency gain; in learning, those same tasks are the mechanism through which the schema forms — so offloading them removes the very difficulty that builds competence. The logic is documented most precisely in software development; we argue it extends to any IT specialisation where AI can perform the domain task before the practitioner can do it independently — though that extension is a reasoned claim, not yet measured beyond programming.

A second difference compounds the first: GenAI has no access barrier. Every earlier tool required the user to adapt — to learn a syntax or interface; GenAI runs on natural language, the one every student already has. Because one can simply talk to it in natural language, Harari (2023) warns that this is a technology that has „hacked the operating system of human civilisation". Needing no learning curve, it can be deployed from day one — and because it responds like a knowledgeable peer, the interaction **feels like reasoning** even though the mechanism beneath it has no understanding, judgment, or accountability. Earlier tools were clearly just tools, but GenAI feels more like a collaborator. This makes its mental impact hard to notice at first, until you actually put it to the test.

Wang et al. (2025) show why a learner's schema completeness — how complete their mental template for a task is — affects the outcome. Practitioners with complete schemas use AI output as raw material, copying selectively and rewriting; those with incomplete schemas accept and reformat it without independent resolution. The same tool produces opposite developmental outcomes — and at the formation stage, the schema is incomplete by definition. Fatigue worsens it, pushing tired students toward delegation at exactly the late-night, end-of-term moments of highest formation value (Kahneman, 2011). Together — schema-building tasks offloaded, no barrier to use, and fatigue-driven delegation — these are what make the cognitive risk cascade possible, which Section 3 traces in detail.

2 Approach and Theoretical Background

This paper is a conceptual synthesis, not a new empirical study: it presents no original data. Its objective is twofold. The first is to integrate established cognitive-science accounts of skill formation with recent empirical findings on AI-assisted learning, tracing a single mechanism: how AI use at the formation stage, without metacognitive discipline, produces never-skilling. The second is to derive from that mechanism a response — hold–rethink–continue, a teachable metacognitive posture operationalised through four student-facing practices.

The synthesis draws on three bodies of evidence. The first is cognitive-science learning theory — cognitive load (Sweller, 1988), deliberate practice (Ericsson et al., 1993), retrieval practice (Dunlosky et al., 2013), self-regulated learning (Zimmerman, 2002; Panadero, 2017), and scaffolding (Vygotsky, 1978; Wood et al., 1976) — which establishes the conditions under which technical competence forms. The second is convergent empirical evidence: three independent randomised controlled trials across distinct domains — coding (Shen & Tamkin, 2026), mathematics (Bastani et al., 2025), and writing (Fan et al., 2025) — that isolate the same mechanism of cognitive bypass. The third is labour-market data on AI adoption in computing (World Economic Forum, 2025; Gartner, 2025; Stack Overflow, 2025). The argument's strength rests on triangulation: established learning science defines what competence formation requires, and the convergent trials demonstrate that current AI use systematically removes those conditions.

These sources were chosen purposively; this is not a systematic review, and the choice is deliberate. The three trials are not the only studies of AI in higher education. Each was chosen because it isolates the same mechanism — a measured gap between assisted output and independent competence — under a randomised design, and because together they span three different domains: coding, mathematics, and writing. That spread is the point. A single mechanism that appears in settings this far apart is hard to dismiss as a quirk of one task. Studies that only describe how much AI is adopted, or how students feel about it, were not drawn on for the mechanism: they cannot speak to whether competence forms.

A note on the status of this evidence, since the field is young and moving quickly. Where a study has reached its journal version, we cite that: the mathematics trial is published in the Proceedings of the National Academy of Sciences (Bastani et al., 2025) and the writing trial in the British Journal of Educational Technology (Fan et al., 2025). The coding trial (Shen & Tamkin, 2026) is still a preprint; we draw on it because it is the closest evidence to IT practice and because the case rests on the agreement of all three trials, not on any single one. The divergence result (Caosun & Aral, 2026) is a formal model, also a preprint, and we treat its prediction as a modelled projection rather than a measured outcome.

The scope is bounded to the IT-education formation stage, where professional schemas and judgment are established; it does not address practising professionals, for whom the dynamics differ; their concern is deskilling (Table 1.).

Table 1. Deskilling vs. never-skilling

Dimension	Deskilling	Never-skilling
Timing	After capability has formed	Before capability has formed
Who	A practising professional	A student at the formation stage
Mechanism	A formed capability erodes through disuse or automation	A capability never forms because AI performs the formative task
Detection	Loss shows against prior performance	No baseline to compare against — invisible until tested independently
Reversibility	Recoverable — retraining restores a known state	Built from scratch; no prior competence to restore
Credential	Reduced mastery behind a still-valid record	A credential without the capability it implies

The two differ in what is at stake, and in the remedy. Deskilling can end a career and is answered by re-skilling. Never-skilling can prevent one from beginning a career; because there are no pre-educational IT skills, it is answered not by recovery but by formation — the hold–rethink–continue posture and the four student practices developed in Section 5.

The paper takes the perspective of the IT student — the individual bearing the career-formation risk — and the student-facing practices that mitigate it. The complementary institutional response — assessment redesign, curriculum, and scaffolding configuration — is developed in a companion paper presented in parallel. Each paper stands independently; together they address the same phenomenon from the learner's and the educator's standpoints.

Because this risk could affect so many students, we think it is better to act on the evidence we have than to wait for proof of every detail. The evidence is indirect and early; we set out its specific limitations, and the research agenda they imply, in the Discussion (Section 6).

3 The Cognitive Risk Cascade

Never-skilling is not a single failure but the possible endpoint of a cognitive risk cascade. Each stage may appear ordinary in isolation — a reasonable shortcut, a missed check, a fluent answer, a passing grade — but together they can carry a student from assisted performance to unformed independent capability without any clear alarm point. Five stages unfold within the individual learner; a sixth then shifts the risk to cohort level.

These six stages build on one another; they are not separate signs to tick off. Each stage is what makes the next one possible: cognitive offloading is what lets

self-monitoring lapse; the monitoring lapse is what lets the illusion of capability form; and that illusion is what lets never-skilling pass unnoticed. Remove an earlier stage and the later ones never arrive — which is why the sequence is a cascade, not a set of independent snapshots.

Stage 1: **Cognitive offloading**

The cascade begins with cognitive offloading: the transfer of cognitive work to an external aid in order to reduce immediate mental demand (Risko & Gilbert, 2016). At the formation stage, generative AI can offload problem decomposition, pattern recognition, error diagnosis, and solution generation before the student has formed a hypothesis or attempted a solution path. Because schemas form through the effortful processing of working a problem through (Sweller, 1988), what is offloaded is precisely what would have built the competence. The tendency runs highest in exactly this age group (Gerlich, 2025), and the choice is not naïve but structurally rational: near-term gains can outrun a distant, discounted skill cost (Caosun & Aral, 2026).

Stage 2: **Self-monitoring lapse**

Offloading turns into risk when self-monitoring lapses — as a habit, not a decision. AI inverts the retrieval sequence that consolidates learning, presenting the answer before the student has tried to produce it (Dunlosky et al., 2013). Fan et al. (2025) measured the result: AI assistance raised output scores while producing no gain in knowledge or transfer — a closed generate-submit loop in which "do I understand this?" is never asked.

Stage 3: **Illusion of capability**

Once monitoring fades, the student loses the internal signal that would flag the gap between performing and understanding. The code compiles, the grade holds, and the self-assessment stays intact — while the capability beneath it has not formed. Fluency in the moment is mistaken for durable mastery (Koriat & Bjork, 2005), and Prather et al. (2023) name the computing form of this the illusion of capability. It is self-concealing by construction: the schema needed to notice the gap is the one the gap has prevented from forming.

Stage 4: **Never-skilling**

As Table 1. shows never-skilling is structurally different from deskilling. Deskilling assumes a skill fades over time, such as when a doctor relies on AI for diagnoses and their own judgment weakens, or when a pilot's flying skills decline after years of using autopilot. But a student cannot lose a skill they never developed in the first place. They simply fail to build it because AI is handling the very work through which it would have formed (Macnamara et al., 2024). The missing ingredient is productive struggle: difficulty, error, and independent resolution are not obstacles to skill but its mechanism (Ericsson et al., 1993). When AI removes all three too early, it can remove the work through which durable competence forms.

Stage 5: **Substitutive AI use**

In the end state, the student stops forming independent judgment and transfers the function to the tool. Scaffolding is meant to fade as the learner grows (Vygotsky, 1978; Wood et al., 1976). A support that never withdraws becomes a dependency, and AI is structurally incapable of withdrawing on its own: it answers every query with the same completeness regardless of who is asking. The output is then accepted rather than evaluated, indistinguishable from competence until the tool is taken away (Shaw & Nave, 2025).

Stage 6: **Cohort divergence**

The first five stages describe what happens inside one student. The sixth changes scale — and this is where the stakes become structural. Caosun and Aral (2026) show that when AI output no longer depends on the user's judgment, a single cohort using a single tool divides in two along the line of prior skill. Table 2. sets out the two trajectories side by side.

Table 2. Divergent AI-use trajectories by prior foundational capability

Formation principle	**Student with foundational capability**	**Student without foundational capability**
Prior schema condition	Arrives with a usable foundational schema	Arrives before the relevant schema has formed
Evaluation of AI output	Filters AI output through existing judgment	Accepts AI output without reliable evaluation
Cognitive engagement	Uses AI output as material for refinement	Treats AI output as an endpoint
Role of AI	AI use becomes complementary	AI use becomes substitutive
Formation mechanism	Capabilities remain in active use	Capability formation is bypassed
Developmental trajectory	Capability accumulates	Never-skilling trajectory begins
Likely outcome	Continued development of independent competence	Risk of a widening capability gap

The same tool, in the same classroom, produces both trajectories at once. For IT students the dividing line falls early: the foundational capabilities for much of the computing curriculum are still forming in the first year or two, so AI adoption at that point is not neutral. It becomes risky not because the tool is harmful in itself, but because the schema that would make its use complementary has not yet formed.

Read forward, the cascade is a warning. Read backward, it is a map: every stage is a point at which engagement, if preserved, can interrupt the cascade. The same model that predicts the widening gap also identifies who escapes it — the students whose use of AI stays complementary rather than substitutive. Whether that is teachable, and what it looks like in practice, is what the rest of this paper sets out to show.

4 Convergent Evidence

No study yet tests this on IT students directly — as Section 2 noted, the evidence is indirect. What it has instead is convergence: three independent randomised controlled trials, in three different settings, isolating the same effect. The closest in domain is a coding-learning task, though its participants were experienced working programmers rather than students. The furthest — school mathematics — matters as much as the closest: when the same effect appears that far from computing, and a single design change switches it off, it is hard to read as an artefact of one setting. Writing sits in between. None is a perfect match for the IT student; taken together they are difficult to dismiss.

4.1 Coding — Shen and Tamkin (2026)

The study used a randomised controlled design: 52 professional and freelance programmers learned a new Python library either with or without AI coding assistance. Both groups produced working code during the session, and the AI-assisted group produced higher-quality immediate output. Afterwards, both sat an independent knowledge assessment with no AI.

On that assessment the AI-assisted group scored about 17 percentage points lower than the control group, with no offsetting productivity gain during the session. The gap was largest in debugging — the skill that most depends on forming hypotheses about failure states, which develops only by meeting and resolving real failures. The unaided group encountered more errors while learning, and higher error-encounter went with better mastery: the productive struggle was not the obstacle but the instruction.

The most useful finding was that the outcome depended on how participants used AI, not whether. Six engagement patterns emerged; the three that kept the participant thinking after the AI answered scored markedly higher than the three that delegated the work. The subgroups are small, so the typology is qualitative rather than precise — but the direction is the point: the performance gap is not between AI users and non-users, but between those who engaged with the output and those who accepted it. That distinction is teachable.

4.2 Mathematics — Bastani et al. (2025)

The effect travels far from coding. In a pre-registered field trial with nearly a thousand school students (grades 9–11), unrestricted GPT-4 access improved performance on AI-assisted practice but left students about 17% worse on unassisted end-of-session exams than controls. The cause was cognitive bypass, not AI error: most requests asked directly for answers rather than method, and the tool's own mistakes were not what predicted exam failure.

The decisive detail is the guardrail. A second condition — a tutor mode giving hints and scaffolding instead of answers — removed the learning penalty entirely while keeping the practice gain. The same tool, redirected from answering to prompting, stopped doing harm: a design that preserves engagement preserves learning.

4.3 Writing — Fan et al. (2025)

The same split appears in writing. In a study of 117 learners on a multi-stage writing-and-revision task, AI assistance raised the quality of the written output but produced no gain in knowledge and none in transfer to a new task: the product improved while the learning did not. Fan and colleagues name the mechanism metacognitive laziness — when the tool handles the monitoring, evaluating, and checking, the learner stops doing them. It is the clearest statement of the pattern running through all three studies: better output is not better learning.

Beyond the three trials, the same shape recurs. Schema-incomplete students engage only at the surface, with no sign of the structure forming (Wang et al., 2025); those who cannot judge AI output accumulate errors through steps that each inherit the last (Reihanian et al., 2024). The signal is already visible in the profession these students are training for: 20% of experienced developers report declining confidence in their own problem-solving, and roughly one in six regularly meet AI-generated code they cannot understand (Stack Overflow, 2025).

5 Student-Facing Practices

By now the case runs in both directions. The cascade is real and measurable: AI use at the formation stage, left unstructured, runs from offloading to never-skilling and leaves a capability gap the student cannot see. While a significant performance gap exists between students who use AI and students who avoid it, a critical secondary gap divides AI users themselves: it lies between those who keep thinking after the tool answers and those who passively accept its output and move on. That difference decides which side of the divergent-trajectory line a student lands on, and it is teachable.

This paper's response is a single posture that makes that engagement deliberate, and it is the paper's central contribution: hold, rethink, continue.

Hold is a short pause before accepting AI output — long enough to ask: "do I actually understand what this is, and why should I accept it?". **Rethink** is reconstructing the reasoning behind the output: why does this solution work, what assumptions sit beneath it, how would it change if the input, the constraints, or the scale were different — and what could go wrong without your noticing? This is not a checklist; it is the habit of stepping back from the result to inspect your own understanding — metacognition in its most practical form. **Continue** is re-engaging the tool with a sharper question: rather than asking for the answer, ask for a hint, a contrasting case, or why an approach fails under a particular constraint, then return better able to judge the response.

Hold–rethink–continue is not a new theory but a way of putting an established one to work. Zimmerman's (2002) model of self-regulated learning runs in three cyclical phases — forethought (goal-setting and planning), performance (monitoring your own work as you go), and self-reflection (evaluating the result and adjusting) — phases Panadero (2017) finds consistent across the major learning-science traditions. Generative AI quietly switches each one off: forethought is skipped when the answer is already there, performance monitoring collapses into accepting the output, and self-reflection never happens when you take the tool's verdict instead of forming your own. Hold–rethink–continue is the posture that switches them back on. The label, and its application to the single decision of when to lean on AI and when to withhold, are this paper's; the underlying cycle is Zimmerman's (2002; Panadero, 2017).

This posture is not idiosyncratic. Reviewing thirty-nine studies of AI and critical thinking, Nasr et al. (2026) independently arrive at a similar structured, human-in-the-loop sequence for keeping the learner in control — convergent derivation from a different literature, not external validation of this one.

The point is not to use AI less, but to use it as the starting point of thinking rather than the end. Four practices put the posture into action, each grounded in the evidence above.

5.1 The Explanation Test

Before submitting any AI-generated output, close the tool and explain the solution's logic in three sentences — not what it does, but why it is built this way. If the explanation comes, the schema is intact and the AI was an accelerant; if it does not, you have accepted a solution you do not understand. It costs under two minutes, and it maps directly onto Shen and Tamkin's (2026) high-scoring patterns, which all required engaging with **why** the output works, not only whether it does.

5.2 Deliberate Non-Use

Pick specific task types — debugging, first-pass decomposition, algorithm design from scratch — and practise them without AI, as a regular discipline rather than a moral stance. Shen and Tamkin's (2026) control group, working unassisted, scored higher on the independent test; the errors they hit were the formation mechanism. Deliberate non-use recreates that condition on purpose, early, before delegation becomes the default — and Bastani et al.'s (2025) guardrail (hints, not answers) shows the same at n≈1,000: the design that preserved the struggle preserved the learning.

5.3 Interaction Pattern Awareness

What separates the high-scoring patterns from delegation is not effort but the direction of attention **after** the AI responds. Conceptual inquiry means asking why — why this algorithm, why this structure, what would make it fail, what alternative was passed over and why. It turns the output from an endpoint into a starting point — and it has to be decided before a session begins, not improvised within it.

5.4 Error-Awareness Practice

Being told you are responsible for AI output does not make you check it: Kupfer et al. (2023) found that participants reminded of their legal accountability changed no verification behaviour, while those told concretely how the tool fails scrutinised far more (d=0.55–0.67). What activates scrutiny is specific knowledge of failure modes. In computing, four recur: hallucinated API calls that do not exist until runtime; the right algorithm solving the wrong problem because the task was under-decomposed; logic-broken code that still passes style and security checks; and confident wrong output on unfamiliar libraries. Those are the conditions to verify against — generic reminders will not produce the check; knowing how the tool breaks will.

6 Discussion

This paper set out to integrate established cognitive-science accounts of skill formation with recent empirical findings on AI-assisted learning into a single mechanism — how generative AI use at the formation stage, absent metacognitive discipline, produces never-skilling — and, from that mechanism, to derive a teachable response. Its contribution is therefore neither a new empirical result nor a literature survey but a synthesis. That synthesis resolves to a single principle: hold–rethink–continue, a metacognitive posture that operationalises self-regulated learning (Zimmerman, 2002; Panadero, 2017) as a concrete decision procedure for when to lean on AI and when to withhold. The principle is earned by the cascade that

precedes it: the cognitive risk cascade (Section 3) reframes a set of separately documented findings as one connected chain. Its evidential weight rests on triangulation rather than on any single study. Established learning science specifies the conditions competence formation requires; three independent randomised controlled trials — across coding (Shen and Tamkin, 2026), mathematics (Bastani et al., 2025), and writing (Fan et al., 2025) — demonstrate that current AI use systematically removes them. Two implications follow.

Theoretically, never-skilling is structurally distinct from deskilling — **a difference of kind**, not degree (Table 1.). That distinction is what makes the double exposure of IT students — a learning environment and a profession restructured by the same technology at once — a categorically different condition from the single disruption other disciplines face.

Practically, the locus of the response is divided: the four student-facing practices in Section 5 that enact hold–rethink–continue are individually low-cost and immediately actionable, while the structural conditions that make them necessary fall to institutions — the educational response developed in the companion paper presented in parallel. What that synthesis means for the individual student — the party who bears the formation risk but did not create it — deserves to be stated plainly, because it is the hardest part of the double exposure to see from inside.

Part of what makes this condition hard to address is that it is hard to perceive from inside. The student arrives to build a career, not to monitor their own cognition, and by then the tool that will shape both is already present — an ambient layer in the editor and the browser that asks nothing to begin. A student can watch the industry change; they cannot as readily watch the same tool change what they think with, not only what they work on. Earlier technologies entered through the curriculum, slowly enough for teachers to update and students to adjust; this one arrived underneath it, and neither side was prepared. When AI helps a student pass, ship, and earn the credential, it looks like nothing but help, and never exposes the fork it conceals: that the same output can be the start of understanding or its end. The individual student can hardly be held culpable for this: a layer one cannot see is not one that was agreed to, and a fork one was never shown is not a choice that was declined.

That blindness is structural, not personal, which is why the response cannot rest on the student alone. The heavier obligation falls on institutions — to redesign assessment, sequence curriculum, and configure AI so that productive struggle survives. It is a prospective duty, not a charge of negligence: the layer arrived faster than any curriculum cycle was built to answer. How institutions should meet it is the subject of the companion paper presented in parallel. But a structural cause does not dissolve the individual stake — it sharpens it.

The professional bottleneck is shifting from generation to verification. AI produces code from natural language with rapidly improving fidelity; what stays scarce — and is intensifying as a requirement — is the judgment to tell whether that code is correct, secure, efficient, and appropriate for its context. That judgment runs on exactly the schemas that delegation at the formation stage prevents from forming. And the gap is easy to miss: "code runs" and „I understand why this code is correct” look identical until the tool is taken away. Stack Overflow (2025) supplies the professional signal: 66% of experienced developers name "AI solutions that are almost right, but not quite" as their primary frustration. The schema needed to catch „almost right” is the one that forms through the struggle delegation removes. Under the divergent trajectories modelled by Caosun and Aral (2026), the students who reach the formation stage without it do not catch up.

Three limitations bound these claims. First, none of the evidence is longitudinal. The trials measure performance immediately after a single session, not months later, so the widening gap between assisted and unassisted students is a model's projection (Caosun & Aral, 2026), not an outcome followed across a career; whether the deficit persists, or can later be recovered, remains open. Second, almost no study observes IT students directly. Even the closest case (Shen & Tamkin, 2026) used working programmers — professional and freelance, of varying experience — learning an unfamiliar library, not students; the argument therefore rests on the three nearest domains available — coding, school mathematics, and writing (Section 4) — which point the same way but leave the effect size for IT students unmeasured. Third, the practices in Section 5 are calibrated to how current AI fails; as the tools improve, those specific failure modes will shift, and the practices will need periodic revision — though the underlying need to verify AI output will not.

These limitations double as a research agenda. The first calls for a longitudinal study that tracks a single cohort across its formation, testing whether assisted and unassisted performance diverge or converge. The second calls for studies situated in IT education itself, across specialisations beyond programming. The third calls for the practices to be revisited as the tools evolve. None is a reason to defer action: too many students are forming their habits now. The direction is consistent with emerging GenAI-specific design work that treats large language models as scaffolds for critical thinking rather than substitutes for learner judgment (Vendrell & Johnston, 2026).

That **hold–rethink–continue** operationalises established self-regulated learning (Zimmerman, 2002; Panadero, 2017), rather than an ad hoc rule, is what makes it a testable proposition — and what the agenda above is designed to test.

7 Conclusion

The window is open now, and it does not stay open. The four practices in Section 5 are not a stand against AI; they are how you make the invisible fork visible and choose the side that builds the schema. Each one keeps you thinking after the tool answers — which is the whole game.

- **Hold** before accepting output you cannot yet explain.
- **Rethink** the reasoning until you could rebuild it without the tool.
- **Continue** — back to AI with a sharper question, now able to judge what it gives you.

Good equipment helps a trained rider win the race. It cannot make you a rider. AI is the equipment; the formation stage is where you become the professional who can use it to win — instead of the one it leaves behind.

References

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122.

Caosun, M., & Aral, S. (2026). The augmentation trap: AI productivity and the cost of cognitive offloading *[arXiv preprint]*. https://arxiv.org/abs/2604.03501

CC2020 Task Force. (2020). Computing curricula 2020: Paradigms for global computing education. *ACM/IEEE-CS/AAAI*.

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving students' learning with effective learning techniques: Promising directions from cognitive and educational psychology. *Psychological Science in the Public Interest*, 14 (1), 4–58.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100 (3), 363–406.

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56 (2), 489–530.

Gartner. (2025). Gartner top strategic technology trends: Software engineering intelligence platforms and AI code assistants. *Gartner, Inc*.

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15 (1), 6.

Harari, Y. N. (2023, April 28). Yuval Noah Harari argues that AI has hacked the operating system of human civilisation. *The Economist.* https://www.economist.com/by-invitation/2023/04/28/yuval-noah-harari-argues-that-ai-has-hacked-the-operating-system-of-human-civilisation

Kahneman, D. (2011). Thinking, fast and slow. Farrar, *Straus and Giroux*.

Ke, Y., Jin, L., Ong, J. C. L., Thirunavukarasu, A. J., Car, J., Cheung, C. Y., Tham, Y. C., Ting, D. S. W., Ong, M. E. H., Compton, S., Narayan, A., Keane, P. A., Wong, T. Y., Bates, D. W., Tan, P., & Liu, N. (2026). AI-induced never-skilling in medical education. *Nature Medicine*, 32, 1997–2006.

Koriat, A., & Bjork, R. A. (2005). Illusions of competence in monitoring one's knowledge during study. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31 (2), 187–194.

Kupfer, C., Prassl, R., Fleiß, J., Malin, C., Thalmann, S., & Kubicek, B. (2023). Check the box! How to deal with automation bias in AI-based personnel selection. *Frontiers in Psychology*, 14, 1118723.

Macnamara, B. N., Berber, I., Çavuşoğlu, M. C., Krupinski, E. A., Nallapareddy, N., Nelson, N. E., Smith, P. J., Wilson-Delfosse, A. L., & Ray, S. (2024). Does using artificial intelligence assistance accelerate skill decay and hinder skill development without performers' awareness? *Cognitive Research: Principles and Implications*, 9, 46.

Nasr, N. R., Tu, C.-H., Sujo-Montes, L., Haniya, S., Yen, C.-J., & Hossain, S. (2026). Generative artificial intelligence impact on critical thinking in higher education: A comprehensive bibliometric analysis and systematic review. *Journal of Computers in Education*.

O'Brien, K. L., & Leichenko, R. M. (2000). Double exposure: assessing the impacts of climate change within the context of economic globalization. *Global Environmental Change*, 10(3), 221–232.

Panadero, E. (2017). A review of self-regulated learning: Six models and four directions for research. *Frontiers in Psychology*, 8, 422.

Prather, J., Denny, P., Leinonen, J., Becker, B. A., Albluwi, I., Craig, M., Keuning, H., Kiesler, N., Kohn, T., Luxton-Reilly, A., MacNeil, S., Petersen, A., Pettit, R., Reeves, B. N., & Šavelka, J. (2023). The robots are here: Navigating the generative AI revolution in computing education. *In Proceedings of the 2023 Working Group Reports on Innovation*

and Technology in Computer Science Education (pp. 108–159). ACM.

Reihanian, I., Hou, Y., Chen, Y., & Zheng, Y. (2024). A review of generative AI in computer science education: Challenges and opportunities in accuracy, authenticity, and assessment. *In International Conference on Computational Science and Computational Intelligence* (pp. 144-158). Springer Nature Switzerland.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688.

Shaw, J., & Nave, C. S. (2025). Thinking — fast, slow, and artificial: Navigating cognitive biases in the age of AI *[SSRN Working Paper]*. https://ssrn.com/abstract=5126527

Shen, Z., & Tamkin, A. (2026). How AI assistance impacts skill formation *[arXiv preprint]*. https://arxiv.org/abs/2603.04825

Stack Overflow. (2025). Stack Overflow 2025 developer survey. https://survey.stackoverflow.co/2025

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12 (2), 257–285.

Vendrell, M., & Johnston, S.-K. (2026). Scaffolding critical thinking with generative AI: Design principles for integrating large language models in higher education. *Computers and Education: Artificial Intelligence*, 10, 100572.

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

Wang, X., Tang, H., & Yu, Q. (2025). Prior domain knowledge shapes cognitive outsourcing patterns in generative AI use: Evidence from a mixed-methods study. *Acta Psychologica Sinica*, 57 (3), 441–459.

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). Future of jobs report 2025. *World Economic Forum*. https://www.weforum.org/reports/the-future-of-jobs-report-2025/

Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory into Practice*, 41 (2), 64–70.