



Teacher-Mediated AI Regulation: A Developmental Equity Framework for K–12 Metacognitive Learning

Sahil Yousuf

National Institute of Technology Srinagar, India

Received : July 24, 2026	Revised : September 3, 2026	Accepted : September 8, 2026	Online : September 23, 2026
--------------------------	-----------------------------	------------------------------	-----------------------------

Abstract

Personalization without autonomy is not necessarily education. Artificial intelligence (AI) can adapt support to individual learners, yet the same system may prompt planning and reflection in one classroom while supplying answers in another. This conceptual paper addresses that divergence through the Teacher-Mediated AI Regulation (T-MAIR) framework for K–12 education. T-MAIR integrates cyclical self-regulated learning, metacognitive monitoring and control, scaffolding and co-regulation, and digital-equity research with recent school-based evidence on adaptive and generative AI. A structured conceptual synthesis, updated through August 2026, treats empirical studies as mechanism-revealing anchor cases rather than as a pooled effectiveness estimate. Recent meta-analytic evidence reports a moderate overall association between AI-supported self-regulated learning ($g = 0.507$), with stronger effects during performance than during forethought, highlighting a performance–competence divide. T-MAIR proposes that durable metacognitive growth depends on four linked conditions: preservation of learner regulatory opportunities, purposeful teacher mediation, readiness-contingent fading of support, and sufficient school capacity to sustain those practices. Five testable propositions follow: P1, higher Learner Regulatory Opportunity (LRO) predicts stronger transfer; P2, teacher mediation changes the function of identical AI support; P3, optimal scaffolding intensity depends on regulatory readiness and task difficulty; P4, school resources affect outcomes partly through mediation capacity; and P5, assisted performance can mask dependence when independent transfer is not assessed. The paper proposes the LRO index, the Frame-Constrain-Interrogate-Fade (FCIF) protocol, and a multilevel mixed-methods design for empirical testing. The defining criterion for successful AI integration is therefore not assisted task completion, but increasingly independent regulation when AI support is reduced or removed.

Keywords: *Artificial Intelligence, Self-regulated learning, Metacognition, Teacher mediation, Digital equity, Educational equity, Scaffolding theory, K–12 education*

INTRODUCTION

Imagine two classrooms using identical AI tutoring software. In Classroom A, the teacher requires students to articulate a learning goal, make an independent attempt, and then consult the AI for targeted feedback. In Classroom B, the same software provides complete solutions on request. Both classrooms use AI-personalized learning, but only one reliably preserves opportunities to practice regulation. The divergence is therefore not explained by AI capability alone. It is a pedagogical decision about who retains regulatory authority. The educational problem is not whether AI can produce useful outputs, but whether AI-supported activity leaves students better able to plan, monitor, evaluate, and revise their own learning.

Self-regulated learning (SRL) offers a precise lens for that problem. SRL refers to the processes through which learners set goals, select strategies, monitor understanding, manage effort, seek help, and evaluate outcomes. Zimmerman's cyclical model organizes these processes into forethought, performance, and self-reflection, while contemporary reviews emphasize that cognition, motivation, affect, and metacognition are mutually regulating rather than separable components (Zimmerman, 2002; Panadero, 2017). AI can intervene at every phase: it can help formulate goals, adapt task difficulty, provide hints and explanations, monitor progress, and

Copyright Holder:

© Sahil. (2026)

Corresponding author's email: sahilyousuf_2024mscphy007@nitsri.ac.in

This Article is Licensed Under:



prompt reflection. However, the same affordances can also assume the regulatory decisions that learners need to practice.

Recent evidence makes this tension visible. [Xu et al.'s \(2026\)](#) meta-analysis of 35 empirical studies reported a moderate overall effect of AI-supported interventions on SRL ($g = 0.507$), with stronger effects during performance ($g = 0.574$) than during forethought ($g = 0.401$). The authors characterize a persistent performance-competence divide: AI may improve what learners can accomplish more readily than it develops durable regulatory competence. School-based studies and recent K-12 syntheses show similar heterogeneity ([Boulhrir et al., 2026](#)). [Li et al. \(2026\)](#) found higher self-regulation, motivation, enjoyment, and engagement after a 10-week AI self-study intervention with 383 primary-school students. By contrast, [Klar \(2025\)](#) found that brief adaptation guidance doubled secondary students' use of chatbot adaptation prompts but did not yield greater knowledge gains, and [Fütterer et al. \(2026\)](#) found no broad knowledge or strategy advantage for targeted GenAI supports over a standard ChatGPT control. These are not contradictory findings so much as evidence that instructional arrangements determine what AI support actually does.

The literature is also rapidly moving beyond a simple 'AI as tool' model. [Molenaar \(2022\)](#) conceptualizes hybrid human-AI learning technologies as a deliberate division of regulatory work. [Kong and Yang \(2024\)](#) place human-centered guidance at the core of K-12 GenAI-supported SRL. [Agustin \(2026\)](#) advances AI as a co-regulator, whose value depends on productive metacognitive friction and the preservation of learner judgment. Most recently, [Salah et al. \(2026\)](#) describe a regulation paradox in higher education: systems intended to support SRL can create regulatory displacement when they assume planning, monitoring, or evaluation that learners should still practice. These developments sharpen, rather than remove, the K-12 gap. What remains under-specified is how teacher mediation, developmental readiness, and school resource conditions jointly determine the allocation and return of regulatory responsibility across elementary and secondary classrooms. This teacher-centered interpretation closely aligns with [Orhani \(2023\)](#), who argues that educational robots are better understood as classroom assistance than as replacements for teachers.

Teacher mediation is central because identical AI output can serve as a shortcut, a model, a feedback cue, or an object of critique, depending on the surrounding activity. Teachers can require an independent attempt before help, constrain the form of assistance, ask students to justify whether an AI response is trustworthy, and deliberately withdraw support during transfer. This role is not reducible to generic 'teacher presence'. It involves concrete moves that connect AI use to phases of SRL and determine who retains decision authority. A primary-school co-regulation study illustrates the nuance: [Horvers et al. \(2024\)](#) found more practice and higher accuracy under a goal-setting and monitoring intervention, yet the control group showed stronger learning gains on the most difficult fraction topic. More support is therefore not automatically better; it must be calibrated to the task and learner readiness. [Dejacto et al. \(2023\)](#) likewise emphasize teacher adaptation, classroom management, and ICT literacy when responding to digitally immersed learners.

Equity introduces a second mechanism. Digital inequality extends beyond device access. Students differ in connectivity continuity, AI literacy, teacher support, class size, language resources, and opportunities to use technology for intellectually demanding rather than merely efficient work ([Hargittai, 2010](#); [Warschauer & Matuchniak, 2010](#)). A well-resourced school can still normalize answer-first AI use, while a resource-constrained school may establish high-quality regulatory routines for limited access. Treating resources and mediation as interacting conditions allows equity to be analyzed as a pathway shaping learning opportunities rather than as a demographic control added after average effects are estimated.

Despite growing attention to AI-supported self-regulated learning (SRL), three limitations remain significant for K–12 theory and practice. First, existing human-centered frameworks often identify teacher involvement as important without specifying observable teacher moves that allocate regulatory decisions between the learner and AI. Second, support is rarely described as a readiness-contingent process that should be intensified, maintained, or faded based on evidence of learner competence. Third, equity is often framed as access or demographic difference rather than as a pathway through which school resources shape the quality and continuity of teacher mediation. These limitations matter especially in K–12 settings because regulatory habits are still developing, developmental differences are pronounced, and schools vary substantially in instructional and technological capacity.

This paper contributes four linked constructs. First, it proposes the Learner Regulatory Opportunity (LRO) index as an observable indicator of decision ownership. Second, it specifies the Frame-Constrain-Interrogate-Fade (FCIF) sequence as a replicable teacher-mediated routine. Third, it models developmental fading as contingent on regulatory readiness rather than on age alone. Fourth, it positions school resources as a mechanism that shapes mediation quality and continuity rather than as an access-only variable. The resulting framework is deliberately falsifiable: its claims can be weakened if LRO does not predict transfer, teacher mediation does not alter AI effects, readiness does not explain support intensity, or resource conditions do not operate through mediation capacity.

The paper proceeds as follows. Section 2 establishes the theoretical foundations and situates T-MAIR within current AI–SRL frameworks. Section 3 explains the conceptual-synthesis method and its analytic safeguards. Section 4 presents five findings and propositions, then translates them into the T-MAIR framework and classroom protocol. Section 5 proposes an empirical test. Sections 6 and 7 present implications, limitations, and research priorities. The final section consolidates the propositions and operational indicators.

The literature has evolved from distinguishing metacognitive monitoring and control to modeling SRL as the coordinated regulation of goals, strategies, motivation, and context, and more recently to examining hybrid human–AI regulation. [Winne and Hadwin \(1998\)](#) emphasize recursive task interpretation, goal setting, enactment, monitoring, and adaptation; [Pintrich \(2000\)](#) emphasizes cognition, motivation, behavior, and context; and [Nelson and Narens \(1990\)](#) distinguish monitoring from the control decisions that follow. Together, these traditions support T-MAIR's claim that AI can assist learning without preserving the learner's opportunity to make the regulatory decision.

Two instructional principles sharpen the framework. [Bjork \(1994\)](#) offers a theoretical counterweight to AI systems that remove all difficulty, whereas [Kirschner et al. \(2006\)](#) show that insufficient guidance can overload novices and that guidance becomes less necessary as prior knowledge grows. T-MAIR therefore rejects both answer-first automation and indiscriminate minimal guidance: support should preserve the regulatory work the learner must eventually perform.

LITERATURE REVIEW

Table 1. Theoretical foundations underlying the T-MAIR framework

Model	Core construct	AI-relevant implication	T-MAIR adaptation
Zimmerman (2002)	Forethought–performance–reflection	AI can support different SRL phases unevenly	Fading protects phase-specific learner decisions
Winne &	Recursive task	AI changes task	LRO codes who

Hadwin (1998)	understanding, goals, enactment, adaptation	conditions and available strategies	interpret and control pivotal decisions
Pintrich (2000)	Cognition, motivation, behavior, context	Regulation depends on learner and context	School resource context becomes a mediation pathway
Nelson & Narens (1990)	Monitoring–control loop	AI may perform monitoring or control that learners need to practice	Interrogation requires learner judgment before acceptance
Scaffolding/co-regulation	Contingent support and shared regulation	Unlimited help can become over scaffolding	FCIF calibrates and fades support

Metacognition and self-regulated learning

Metacognition concerns knowledge of cognition and the regulation of cognitive activity. Flavell's early formulation distinguished awareness of one's cognitive processes from the monitoring and control that guide them. [Schraw and Dennison \(1994\)](#) later operationalized metacognitive awareness along the knowledge-of-cognition and regulation-of-cognition dimensions. Conceptual reviews likewise emphasize that monitoring and control recur throughout learning rather than appearing only at the end of a task ([Greene & Azevedo, 2007](#); [Veenman et al., 2006](#)). In school learning, these processes are embedded within broader SRL systems that also include motivation, behavior, strategy selection, and resource management. [Pintrich and De Groot's \(1990\)](#) study of seventh graders demonstrated that self-efficacy, intrinsic value, and self-regulation were meaningfully associated with classroom performance, reinforcing the view that successful learning is not reducible to cognitive ability alone.

The field has shifted from asking whether technology can deliver instruction to asking how technology participates in regulation. Adaptive tutoring emphasized personalization and feedback; generative AI expanded planning, explanation, critique, and reflection; the emerging regulatory perspective asks who performs these functions and what happens when support disappears. T-MAIR narrows this perspective to school ecology by treating teacher mediation, readiness, resource continuity, and independent transfer as interacting mechanisms.

[Zimmerman's \(2002\)](#) cyclical model is particularly well suited to AI-supported learning because it highlights temporal phases. Forethought includes task analysis, goal setting, and motivational beliefs. Performance includes strategy use, self-observation, help seeking, and attention control. Self-reflection includes self-evaluation, causal attribution, and decisions about subsequent action. A key implication is that an intervention can improve one phase while leaving others unchanged. This helps explain why some AI studies report stronger changes in engagement or motivational value than in durable knowledge or strategy transfer. [Panadero's \(2017\)](#) comparative review similarly emphasizes that SRL is a system of interacting processes rather than a single trait.

The school intervention literature establishes that SRL can be taught. Meta-analytic evidence links SRL to academic achievement across school years and shows that strategy instruction can improve performance ([Dent & Koenka, 2016](#); [Donker et al., 2014](#)). [Dignath and Büttner's \(2008\)](#) meta-analysis across primary and secondary education found positive intervention effects on achievement, strategy use, and motivation, and showed that design characteristics matter. [Theobald's \(2021\)](#) later meta-analysis in higher education also found that feedback, cooperative arrangements, and metacognitive theoretical grounding moderated training effects. Although the age group differs, the general lesson remains relevant: instructional architecture determines whether regulatory skills are practiced explicitly enough to improve.

AI as a regulator of learning

AI can intervene at several points in the SRL cycle. Adaptive systems can diagnose response patterns and select subsequent tasks. Intelligent tutors can provide hints, prompts, and feedback. Learning analytics can make progress visible. Generative AI adds dialogic flexibility: it can ask planning questions, generate explanations at different levels, simulate a peer, challenge reasoning, or support self-evaluation. Recent reviews therefore describe substantial potential for AI-supported SRL but also recurring limitations, including inaccuracy, learner dependence, weak AI literacy, and an overconcentration of support in the performance phase (Xia et al., 2026; Nguyen et al., 2026). A systematic review by Guan et al. (2025) likewise found that educational chatbots most often support strategy enactment and monitoring, while goal setting, planning, reflection, and adaptation remain comparatively under-supported. Education-focused discussions of ChatGPT similarly emphasize its potential for personalized interaction alongside challenges related to verification, responsible use, and pedagogy (Kasneci et al., 2023; Lo, 2023; Tlili et al., 2023).

Empirical findings are mixed in ways that are theoretically informative. Fütterer et al. (2026) randomized 371 Grade 7–9 students into generative-AI conditions targeting utility value or cognitive learning strategies, and a standard-ChatGPT control. The targeted utility-value intervention produced a more favorable trajectory in perceived usefulness, but the experimental conditions did not show broad advantages in effort, domain knowledge, or elaboration. The result challenges the assumption that adding an AI scaffold automatically yields transfer. It also suggests that the content and timing of prompts must align with a specific regulatory mechanism.

Klar (2025) examined 106 Grade 9–10 students who used a generative-AI chatbot for research. Adaptation guidance more than doubled the use of adaptation prompts without increasing cognitive load, yet both groups showed similar knowledge gains. This distinction between tool use and learning outcomes is crucial. Learners may become more competent users of AI without becoming more knowledgeable, just as a system may reduce friction without deepening processing. Effective educational design therefore needs to distinguish among interaction success, task performance, metacognitive growth, and independent transfer.

At the elementary level, Li et al. (2026) reported that 383 primary-school students in AI-enabled self-study rooms showed higher post-intervention motivation, self-regulation, enjoyment, and engagement than peers in traditional self-study settings. The study indicates that structured AI environments can support younger learners, but it also highlights a common measurement limitation: aggregate pre- and post-scales provide less information about which interaction sequences produced change. Process-level evidence is needed to determine whether gains reflect improved planning and monitoring, stronger engagement, or simply more supportive task environments.

Human mediation, scaffolding and co-regulation

Scaffolding theory offers a clear way to conceptualize the teacher's role. In classic accounts, effective support is contingent on learner need and is gradually withdrawn as competence grows (Wood et al., 1976). Computer-based scaffolding research extends this principle to digital environments but repeatedly warns that fixed or overly directive support can suppress strategic activity (Azevedo & Hadwin, 2005). AI intensifies this tension because it can deliver assistance instantly and indefinitely. The convenience of unlimited help can become educationally counterproductive if students never experience the productive difficulty required to diagnose confusion, retrieve information, or construct an explanation.

Co-regulation shifts the focus from what the technology does to how responsibility is jointly managed. Hadwin et al. (2011) describe regulation as shifting among self-, co-, and socially shared

forms, depending on the task and context. In a teacher–AI–student triad, the teacher can regulate the regulation: deciding when AI should prompt, when the learner must attempt independently, when an output should be challenged, and when support should fade. [Horvers et al. \(2024\)](#) provide school-level evidence for this logic. In their quasi-experimental Grade 5 study, goal setting and progress monitoring were embedded within an adaptive mathematics environment rather than left to the platform alone. The intervention did not simply add more technology; it changed the regulatory structure surrounding technology use. (See also [Järvelä & Hadwin, 2013](#)).

Feedback research also supports a mediation perspective. [Hattie and Timperley \(2007\)](#) distinguish feedback about the task, process, and self-regulation, noting that feedback is powerful only under specific conditions. AI can generate large volumes of feedback, but volume is not the same as quality. A teacher can convert generic feedback into a metacognitive episode by asking students to explain why a suggestion is appropriate, identify evidence for it, or predict whether the same strategy will work on a new task. Such moves transform feedback from information received into regulation enacted. Recent experimental evidence also shows that explicit metacognitive support can strengthen SRL in GenAI environments, reinforcing the importance of instructional mediation rather than tool exposure alone ([Xu et al., 2025](#)).

Developmental calibration

Development changes the amount and form of external regulation learners require. Metacognitive knowledge and monitoring improve across childhood and adolescence, though development is uneven and strongly shaped by task familiarity. A younger learner may benefit from explicit prompts that name a strategy, model its use, and cue reflection. An older learner may need less prompting but greater challenge to justify decisions, compare sources, and detect when convenience is replacing thought. The principle is not that older students always need less support; rather, support should be calibrated to the regulatory process that is not yet stable. Longitudinal K–12 evidence also makes the developmental issue explicit: [Zhang and Derakhshan \(2025\)](#) examined GPT-4o use in K–12 settings over time, providing a useful basis for treating developmental change and task engagement as longitudinal rather than purely cross-sectional phenomena.

T-MAIR's novelty claim is deliberately specific rather than absolute. Existing frameworks already establish human guidance, autonomy, co-regulation, and responsible AI use matter. T-MAIR adds an operational question: which regulatory decision is retained by the learner, when assistance is introduced, how support is adjusted or faded, and what school conditions make that allocation sustainable? In this sense, the framework is an integration and operationalization contribution rather than a claim to have invented teacher mediation or regulatory displacement.

This distinction can reconcile seemingly contradictory findings. When low-SRL learners receive structured AI support, the tool may externalize useful processes they would otherwise omit. When high-SRL learners receive the same level of prompting, it can constitute overscaffolding, reducing self-initiated strategy use. Recent AI-SRL reviews therefore argue for adaptive support that accounts for prior self-regulation and gradually returns control to the learner ([Xia et al., 2026](#); [Nguyen et al., 2026](#)). A developmental framework should therefore model support as a dynamic allocation of responsibility, not as a fixed platform feature.

Digital inequality as an outcome pathway

Digital-equity research has increasingly moved beyond a binary access/no-access model. Differences in skills, quality of use, social support, and the benefits derived from technology can persist even after device access becomes widespread. This logic is especially important for AI, because productive use requires a combination of literacy, verification practices, domain

knowledge, and pedagogical framing. Long and Magerko (2020) define AI literacy as a set of competencies needed to understand and critically evaluate AI, while Ng et al. (2021) emphasize knowing, using, evaluating, and ethically engaging with AI. These competencies are not distributed uniformly across schools or households. Related systematic-review evidence on AI-mediated language learning also cautions that learner autonomy is difficult to infer when AI supports only task-specific phases rather than the full regulatory cycle (Chang & Sun, 2024).

Teacher capacity is part of the inequality mechanism. In a high-resource context, teachers may receive professional development, have time to review AI outputs, and rely on reliable connectivity that supports planned lesson routines. In a constrained context, the same teacher may use AI mainly as a time-saving answer generator because classes are larger, devices are shared, or connectivity is unstable. Thus, resource context moderates pedagogy, not merely access. This is why equity should be conceptualized as a relationship among material access, human mediation, learner capability, and outcomes. Evidence from Valenzona (2026) similarly shows that the effectiveness of digital and instructional materials depends on how teachers select and integrate them, rather than on availability alone. This is reinforced by Research Synergy Press evidence showing that AI awareness and digital competence among basic-education teachers are meaningfully related, and that AI/ICT training is associated with AI awareness (Manaig et al., 2025).

Positioning T-MAIR among current AI-SRL frameworks

T-MAIR is designed as an extension of, rather than a replacement for, current human-AI regulatory frameworks. Molenaar's (2022) hybrid-intelligence account clarifies that control can shift between human and machine, while Kong and Yang (2024) foreground human-centered guidance in K-12 SRL. Chiu (2024) links GenAI activities to autonomy, competence, and relatedness through self-determination theory. Two 2026 contributions make the regulatory risk more explicit: Agustin (2026) treats AI as a relational co-regulator that should preserve productive metacognitive friction, and Salah et al. (2026) use regulatory displacement and bounded autonomy to analyze agentic AI in higher education. The distinctive role of T-MAIR is narrower and more school-ecological: it specifies teacher mediation as an observable mechanism, makes developmental fading explicit across elementary and secondary schooling, embeds resource continuity and teacher capacity as equity pathways, and defines successful support by independent transfer rather than assisted performance alone. This positioning aligns with a recent systematic review of 73 studies, which identifies personalization, feedback, information search, and problem-solving as major GenAI affordances while calling for designs that more deliberately support the full SRL cycle (Xia et al., 2026).

Table 2. Frameworks and school-based evidence informing T-MAIR

Framework	Primary focus	Human mediation	Developmental calibration	Equity/resource mechanism	T-MAIR relation
Molenaar (2022)	Hybrid human-AI regulation	Shared control/augmentation	Implicit	Not central	T-MAIR operationalizes classroom mediation and fading
Chiu (2024)	GenAI activities through SDT	Teacher-supported design	Limited	Not central	T-MAIR adds regulatory decision authority and transfer

Kong & Yang (2024)	Human-centered K-12 GenAI + SRL	Teacher guidance central	K-12 oriented	Not a primary mechanism	T-MAIR specifies observable mediation moves
Agustin (2026)	AI as relational co-regulator	Relational / interactional	Developmental growth valued	Discussed broadly	T-MAIR locates co-regulation in teacher-AI-student classroom ecology
Salah et al. (2026)	Agentic AI, bounded autonomy	Teacher/institutional governance	Capability-sensitive	Institutional readiness	T-MAIR translates displacement risk to K-12 pedagogy
T-MAIR (this paper)	K-12 metacognitive development	Frame-Constrain-Interrogate-Fade	Explicit readiness-based fading	Resources → capacity → mediation	Independent transfer is the success criterion
Study	Level	Sample/design	AI/support	Verified outcome pattern	T-MAIR implication
Iterbeke et al. (2021)	Grades 8–9	n = 1,177; 32 schools; school-level RCT	Adaptive instruction × elaborated feedback	Program ≈ +0.47 SD knowledge; adaptive features added no extra learning gain	Adaptivity itself does not guarantee added value
Ng et al. (2024)	Secondary 4	n = 74; 3-week comparative study	GenAI SRLbot vs rule-based chatbot + teacher/guides	GenAI condition improved science knowledge, engagement and motivation	AI effects occurred within a guided instructional ecology
Horvers et al. (2024)	Grade 5	n = 114; 4 schools; quasi-experimental	Goal setting + monitoring around adaptive mathematics	More practice and higher accuracy; control stronger on hardest topic	Co-regulation can help process yet still overburden difficult learning
Klar (2025)	Grades 9–10	n = 106; experiment	Adaptation guidance for GenAI chatbot	Adaptation prompts M = 6.0 vs 2.5; similar knowledge gains	Brief guidance changes use behavior, not necessarily learning
Zhang &	High	n = 165;	GPT-4o-	Treatment	Longer

Derakhshan (2025)	school	semester longitudinal study	supported learning vs comparison	showed stronger trajectories in SRL and task engagement	exposure may matter, but transfer and mediation remain open
Fütterer et al. (2026)	Grades 7–9	n = 371; preregistered RCT	GenAI utility-value vs strategy prompts vs standard ChatGPT	No broad advantage over control for effort, knowledge or elaboration	Theory-informed prompting requires stronger task-process alignment
Li et al. (2026)	Primary school	n = 383; 10-week quasi-experiment	AI self-study rooms vs traditional self-study	Self-regulation $\eta^2 = .238$; motivation, enjoyment, engagement also higher	Strong assisted-environment effects require tests of independent transfer

Taken together, the literature reviewed indicates that the educational value of AI-supported learning depends not only on the availability or sophistication of AI tools but also on how learners are guided to monitor, evaluate, and gradually regulate their use of these tools. The evidence also highlights important gaps in developmental differences, teacher-mediated regulation, scaffolding intensity, resource continuity, and independent transfer when AI support is reduced. These gaps provide the rationale for the conceptual synthesis undertaken in the next section. Section 3 therefore describes the research methodology used to identify, organize, and synthesize relevant evidence and to derive the Teacher-Mediated AI Regulation (T-MAIR) framework and its associated propositions.

RESEARCH METHOD

Reflexivity is part of the analytic standard. Anchor cases were purposively selected to highlight contrasts across age, AI function, mediation, duration, and outcome. This increases theoretical variation but can overrepresent studies that make regulatory mechanisms visible, so claims about prevalence are unwarranted. Future systematic work should preregister searches, use independent coders, report reliability, document how disagreements were resolved, and obtain external expert review.

This paper uses a structured conceptual synthesis and theory-building approach. The purpose is explanatory integration: to connect established SRL and metacognition theory with recent K-12 AI evidence, identify mechanisms that account for heterogeneous outcomes, and derive directly testable propositions. The approach is not presented as a systematic review or meta-analysis; accordingly, no pooled effect is calculated, and the included empirical studies are treated as anchor cases that reveal contrasts in age, AI function, mediation, duration, and outcome measurement.

Evidence identification and selection

Evidence identification followed a reproducible keyword protocol updated through 23 August 2026. Search combinations paired AI terms (artificial intelligence, generative AI, ChatGPT,

adaptive learning, intelligent tutoring) with regulatory terms (self-regulated learning, metacognition, self-regulation, co-regulation) and school terms (K-12, primary, elementary, secondary, high school). Priority was given to peer-reviewed publisher records and DOI-verified articles. Foundational SRL, scaffolding, feedback, and digital-inequality sources were retained if they provided constructs needed to interpret the school evidence. Recent reviews and meta-analyses were used to assess whether the mechanism claims were consistent with the wider field. To make the revision auditable, the evidence base was also checked against publisher and DOI records for bibliographic accuracy, and the revised reference list was reconciled with in-text citations.

Seven school-based empirical studies were selected as anchor cases (Table 2) because they span primary and secondary education, conventional adaptive technologies and generative AI, short and semester-length exposure, and positive, null, and mixed findings. The selection deliberately included disconfirming evidence: adaptive features with no additive learning benefit ([Iterbeke et al., 2021](#)), co-regulation that improved practice behavior but not the most difficult learning outcome ([Horvers et al., 2024](#)), and GenAI guidance that changed interaction behavior without increasing knowledge ([Klar, 2025](#)). This negative-case logic reduces the risk of constructing a framework based solely on favorable AI outcomes.

Concept-development procedure

Concept development proceeded in four passes. First, each anchor case was mapped to the forethought, performance, and reflection processes. Second, the locus of regulatory decision authority was coded as predominantly learner-led, jointly mediated, or system-led. Third, contextual conditions were mapped at the learner (readiness, prior knowledge, AI literacy), classroom (teacher mediation, task design, class size), and school (access continuity, technical support, professional development) levels. Fourth, the emerging propositions were compared with current hybrid-regulation, human-centered, co-regulation, and bounded-autonomy frameworks to identify overlap and genuine differentiation. Claims that could not be distinguished from existing theory were revised rather than presented as novel.

Analytic standard

The analytic standard separates published evidence, theoretical inference, and proposed measurement. Numerical values reported in empirical studies are reproduced only when they can be verified against publisher or indexed source records. T-MAIR pathways are theoretical propositions, not causal estimates from the present paper. The Learner Regulatory Opportunity (LRO) index introduced below is likewise a proposed coding metric, not a validated instrument. This separation is important in a fast-moving field where conceptual language can become outdated quickly; for example, 'regulatory displacement' is treated here as an established 2026 concept rather than as a new term ([Salah et al., 2026](#)).

FINDINGS AND DISCUSSION

Regulatory preservation versus regulatory displacement

Across the evidence base, the most informative distinction is not AI versus no AI but whether the learning episode preserves regulatory work for the student. Recent theory describes the corresponding risk as regulatory displacement: AI assumes planning, monitoring, evaluation, or strategic decisions before the learner has an opportunity to perform them ([Salah et al., 2026](#)). [Agustin \(2026\)](#) similarly argues that productive metacognitive friction should not be smoothed away merely because automation can do so. T-MAIR adopts this problem but reframes it as a school-

level question: at each decision point, who must notice the problem, generate a strategy, judge an explanation, and decide what to do next? Burden reduction is educationally useful when it removes irrelevant friction; displacement becomes risky when it removes practice of the target regulatory capability itself.

The empirical pattern supports a mechanism-sensitive account. [Fütterer et al. \(2026\)](#) found that the trajectory of perceived utility value differed across targeted GenAI conditions, but the utility-value intervention did not outperform the standard-ChatGPT control on that outcome, and neither targeted intervention produced broad advantages in effort, tested domain knowledge, or elaboration. [Klar \(2025\)](#), meanwhile, showed that very brief adaptation guidance changed how students prompted a chatbot without changing knowledge gain. These findings imply that better interaction and better learning are related but not equivalent. The regulatory process that the intervention is intended to develop must be measured directly and, ideally, again when AI is unavailable.

To make this claim observable, T-MAIR proposes a provisional Learner Regulatory Opportunity (LRO) index. A regulatory decision point is any moment when a learner could set or revise a goal, choose a strategy, judge understanding, request or defer help, verify a claim, or evaluate an outcome. LRO is calculated as the proportion of coded decision points at which the learner generates or explicitly justifies the decision before a teacher or AI system completes it. A high LRO does not mean 'less support'; a teacher may provide substantial scaffolding while still requiring the learner to make and justify the pivotal decision. The index therefore separates support intensity from decision ownership and can be estimated from observations, think-alouds, or platform traces. Its reliability and predictive validity must be established empirically before it is used as an outcome measure.

Proposition 1: Holding content and time-on-task constant, AI-supported learning episodes with greater Learner Regulatory Opportunity will yield stronger growth in metacognitive regulation and in delayed no-AI transfer than episodes in which AI makes the same regulatory decisions for the learner.

Teacher mediation changes the function of the same AI

AI features do not have fixed educational effects. A generated explanation can serve as a shortcut, a worked example, a misconception detector, or a stimulus for critique, depending on the teacher's orchestration. In [Ng et al. \(2024\)](#), a GenAI SRL chatbot was embedded in a three-week science experience in which both groups were taught by the same experienced teacher and received guiding materials; the positive pattern cannot be interpreted as a context-free chatbot effect. [Horvers et al. \(2024\)](#) likewise show that explicit goal setting and monitoring in adaptive mathematics altered practice behavior but did not improve every learning outcome. The mechanism is therefore the interaction among tool affordance, task demand, and human regulation.

The teacher's role is especially important because AI feedback lacks full knowledge of classroom intentions, learner history, and the pedagogical value of withholding an answer. T-MAIR specifies four observable mediation moves. Frame establishes what the learner must ultimately do independently. Constrain protects an initial space for learner-generated work. Interrogate turns AI output into a claim that must be checked, compared, or explained. Fade returns responsibility through reduced prompting and no-AI transfer. These moves convert 'teacher presence' from a background variable into a manipulable component that can be observed for fidelity and tested as a moderator or mediator.

Proposition 2 follows: the effect of AI support on metacognition will be moderated by the quality of teacher mediation, with the strongest gains occurring when teachers explicitly connect AI interactions to goal setting, monitoring, justification, and reflection. Importantly, the proposition

predicts not merely an additive teacher effect but a change in the meaning of AI use itself.

The optimal amount of support is developmental and learner-specific

Elementary and secondary learners should not be treated as interchangeable users, but age should not be treated as destiny. Younger learners often require more explicit models of goal quality, monitoring cues, and help-seeking, whereas older learners can usually manage greater choice and critique. Within each grade band, however, baseline SRL, prior knowledge, AI literacy, and task complexity produce substantial variation. The relevant construct is regulatory readiness: the learner's current capacity to make productive decisions with decreasing external control. Developmental stage informs readiness but does not replace its measurement.

Figure 3 illustrates this as a readiness-based fading principle. The curves are conceptual, not normative age trajectories. External regulation should decrease only when evidence shows that the learner can reliably make the corresponding decision. Conversely, support may need to increase temporarily when task complexity rises. Fading is therefore contingent and reversible, consistent with scaffolding theory rather than a one-way withdrawal schedule.

[Horvers et al. \(2024\)](#) present an instructive boundary case. Grade 5 students in the co-regulation condition solved more practice problems and achieved higher practice accuracy, yet learning gains on complex fractions favored the control condition. This pattern cautions against assuming that a regulatory prompt is beneficial merely because it increases visible activity. On difficult tasks, additional monitoring requirements may compete with domain processing or impose premature self-regulatory demands. T-MAIR therefore predicts an interaction among support intensity, task difficulty, and readiness rather than a universal main effect of increased scaffolding.

Proposition 3: The relationship between AI scaffolding intensity and metacognitive growth will be curvilinear, with under-scaffolding harming low-readiness learners and over-scaffolding reducing self-initiated regulation among higher-readiness learners. Developmental stage and baseline SRL should shift where the optimum lies.

This finding advances current AI–SRL work by making transfer—not assisted performance—the decisive outcome. It challenges the assumption that higher task scores under AI support are sufficient evidence of learning. Its practical implication is that every intervention should include at least one structurally comparable task in which the learner must regulate without the same AI assistance. Without that counterfactual, personalization and dependency can be difficult to distinguish.

Equity depends on mediated opportunity, not access alone

The equity analysis identifies four zones (Figure 4). Low-resource/low-mediation settings face the greatest risk because students encounter both discontinuous access and a weak regulatory structure. High-mediation/low-resource settings can partially compensate through shared-device routines, teacher modeling, and carefully timed use, although scarcity still limits exposure and continuity. High-resource/low-mediation settings reveal a less visible inequality: students may enjoy reliable access but mostly practice answer acquisition rather than independent regulation. High-resource/high-mediation settings offer the strongest conditions for calibrated support and fading. The matrix therefore treats equity as the distribution of high-quality regulatory opportunity, not merely the distribution of devices.

This framework reframes the digital divide. Device counts and bandwidth remain necessary indicators, but they do not capture the pedagogical opportunities students receive. Two schools may report identical AI adoption rates yet differ dramatically in whether students are taught to

verify outputs, articulate goals, or work independently after receiving assistance. These process differences are likely to accumulate over time. Students who already possess strong domain knowledge and metacognitive skills can use AI more productively, whereas learners with fewer supports may accept outputs uncritically or use AI primarily for completion.

Table 3. Frame-Constrain-Interrogate-Fade (FCIF) teacher-mediation protocol

FCIF move	Teacher action	Learner regulatory opportunity	Example
Frame	State the independent target and success criterion before AI use	Goal setting and task interpretation	"First state what a strong answer must accomplish."
Constrain	Require an initial attempt or plan before substantive assistance	Strategy generation and help-seeking choice	"Show your attempt before requesting a solution."
Interrogate	Require evidence-based checking of AI output	Monitoring, evaluation, verification	"What evidence supports or contradicts the AI response?"
Fade	Reduce prompts and schedule comparable no-AI transfer	Adaptation and independent regulation	"Complete a similar task without AI and explain your strategy."

Proposition 4: School resource context will indirectly moderate AI effects through teacher capacity, continuity of access, and opportunities for guided practice. After controlling for device access, schools with stronger teacher AI literacy and structured mediation routines should show smaller between-student gaps in metacognitive growth.

Average achievement is an incomplete equity outcome


AI evaluation often prioritizes immediate performance: quiz scores, task completion, response accuracy, or engagement while support is available. These outcomes matter, but they can overstate durable learning when the system provides strategic support. [Xu et al. \(2026\)](#) found a moderate overall effect of AI-supported SRL interventions ($g = 0.507$), with a larger performance-phase effect ($g = 0.574$) than forethought ($g = 0.401$) and explicitly identified a performance-competence divide. For T-MAIR, the decisive test is therefore counterfactual: what can the learner regulate when the AI prompt, recommendation, or explanation is no longer present? Immediate assisted performance and delayed independent competence should be reported as distinct outcomes rather than collapsed.

A stronger measurement model should integrate four layers. First, self-report scales capture metacognitive awareness, self-efficacy, and perceived strategy use. Second, behavioral traces capture help-seeking, time on task, revisions, prompt types, and the timing of assistance. Third, performance tasks assess near- and far-transfer under conditions where AI is unavailable. Fourth, observation and interview data reveal teacher mediation, student reasoning, and contextual constraints that logs cannot infer. This combination addresses a long-standing problem in SRL research: the gap between what learners say they do and what they actually do during tasks. Recent psychometric work further suggests that GenAI-supported SRL is multidimensional, spanning cognitive, metacognitive, motivational, affective, social-behavioral, and environmental regulation ([Wang et al., 2026](#)).

Proposition 5: AI interventions that appear equivalent in immediate achievement will diverge in delayed independent transfer when they differ in the extent of learner-generated monitoring and reflection. This proposition provides a practical criterion for distinguishing personalization from dependency.

T-MAIR also differs from alternative positions. It rejects technological determinism because identical AI affordances can yield different outcomes under different mediation routines. It extends critical AI literacy by treating critique as one regulatory move within a broader allocation of decision ownership. It qualifies minimal-guidance arguments by recognizing that novices may need substantial support while higher-readiness learners may need greater autonomy. It extends questioning-AI approaches by requiring that questioning not permanently replace learner-generated monitoring and evaluation. The common principle is calibrated responsibility.

Boundary conditions should be explicit. T-MAIR is developed primarily for formal K–12 academic learning in classroom settings where teachers can influence AI use. It assumes at least intermittent access to the instructional environment and focuses on regulating learning tasks rather than on clinical, therapeutic, or purely vocational applications. The framework is not intended to imply that all support should eventually disappear: in some contexts, co-regulation and teacher authority may remain legitimate and valuable. Its central prediction concerns transferable decision ownership, not independence as an absolute educational endpoint.



Core principle: support is educationally valuable when it enables regulation without permanently replacing the learner's regulatory work.

Figure 1. T-MAIR multilevel process model linking school resources, developmental readiness, teacher mediation, SRL phases, Learner Regulatory Opportunity, and independent transfer.

The T-MAIR framework

The Teacher-Mediated AI Regulation framework integrates findings across four interacting levels (Figure 2). At the learning-episode level, forethought, performance, and reflection include regulatory decision points. At the interaction level, teachers use Frame-Constrain-Interrogate-Fade moves to allocate decision authority between the learner and AI. At the developmental level, regulatory readiness determines how much external support is appropriate and when it should be faded. At the institutional level, resource continuity and teacher capacity determine whether high-quality mediation can be sustained. The criterion of success is metacognitive growth, expressed as increasingly independent planning, monitoring, evaluation, and transfer.

T-MAIR's contribution is not the general claim that humans should remain involved with AI; current frameworks already establish that principle. Its added value lies in connecting four elements that are usually treated separately in K-12 research: observable teacher mediation, readiness-based fading, school-resource pathways, and independent transfer. This produces a causal map that is empirically testable. If teacher mediation does not moderate outcomes, if LRO does not predict transfer, if developmental calibration adds no explanatory value, or if resource effects do not operate through mediation capacity, the core T-MAIR claims would be weakened. Falsifiability is therefore part of the framework's design rather than an afterthought.

Evidence-informed implementation model

The practical implication is a four-move sequence: Frame, Constrain, Interrogate, and Fade. Frame begins with an independent learning target and success criterion before the tool is opened. Constrain preserves an initial attempt, prediction, plan, or explanation before substantive AI help. Interrogate requires the learner to treat AI output as contestable—by checking evidence, comparing alternatives, identifying uncertainty, or explaining why a response should be accepted. Fade returns responsibility by reducing prompts, limiting AI availability, and scheduling comparable no-AI transfer. The sequence is tool-agnostic and can be implemented with adaptive platforms, tutoring systems, or generative chatbots. The insistence on an initial attempt and later retrieval/transfer is consistent with evidence that effective learning techniques require learners to retrieve, elaborate, and monitor rather than only re-read or receive completed solutions ([Dunlosky et al., 2013](#)).

In elementary classrooms, framing and constraints are often more visible: structured goal cards, predefined help stems, worked verification examples, and short reflection routines. In secondary classrooms, mediation can shift toward student-selected support, justification for why AI was used, and greater responsibility for source checking and strategy choice. These are starting assumptions, not fixed prescriptions. Fidelity should be scored against the move's regulatory function, while the intensity of support is adjusted to individual readiness and task difficulty.

Contributions and theoretical differentiation

The framework makes three theoretical contributions. First, it translates broad calls for 'human-centered' AI into an allocation-of-regulatory-responsibility model: the unit of analysis is the decision point at which a learner, teacher, or system sets a goal, selects a strategy, judges understanding, or evaluates an outcome. Second, it links co-regulation to development by making fading conditional on regulatory readiness rather than on age or exposure alone. Third, it treats equity as a mechanism: school resources shape teacher capacity and the continuity of guided practice, which in turn shape the regulatory opportunities students actually receive. This is more specific than adding socioeconomic status as a control variable after an intervention effect is estimated.

Methodologically, T-MAIR proposes triangulation rather than a single SRL score. The LRO index captures decision ownership; observation captures mediation fidelity; platform traces capture help-seeking, verification, and revision; self-report captures perceived regulation; and delayed no-AI transfer captures whether regulation persists. Practically, the Frame-Constrain-Interrogate-Fade sequence is intentionally minimal. Its value can be tested on existing school platforms without requiring a new AI system, making the framework suitable for pragmatic trials as well as controlled studies.

Implementation constraints and design safeguards

Implementation quality is constrained by teacher time, platform configurability, student

expectations, and infrastructure. A protocol requiring teachers to inspect every prompt is unlikely to scale. T-MAIR therefore favors low-burden routines: one independent-attempt rule, a small set of verification prompts, sampled rather than continuous review of AI interactions, and short transfer tasks that double as formative assessment. When platforms cannot delay answers or expose reasoning steps, teachers can impose activity-level constraints through worksheets, prompt templates, or staged access. In low-connectivity settings, shared-device demonstrations and paired critique can preserve regulatory practice even when one-to-one access is impossible.

A second safeguard concerns trust. Students should neither accept AI output automatically nor reject it performatively. Interrogation should be evidence-directed: What claim is being made? What information supports it? What could be wrong? What alternative strategy would yield the same result? This aligns the framework with human-computer interaction research on calibrated trust while keeping the educational objective distinct: the learner must increasingly perform the checking rather than merely observe the system's confidence display. Finally, schools should treat transfer evidence as a governance signal. If assisted performance improves while independent transfer stagnates, the correct response is not necessarily more AI exposure; it may be a redesign of how support is allocated and faded.

Teacher-Mediated AI Regulation (T-MAIR): Multilevel Process Model





Figure 2. T-MAIR multilevel process model linking school resources, teacher mediation, SRL phases, developmental calibration, and independent transfer.



Grade band is an informative proxy, but baseline SRL, prior knowledge, AI literacy and task difficulty should also determine fading.

Figure 3. Readiness-based developmental fading principle; grade band informs but does not determine support intensity.




Figure 4. Equity matrix: resource continuity × teacher mediation capacity

T-MAIR reframes the central educational question from whether AI can personalize assistance to whether personalization increases learners' capacity to regulate learning. The framework identifies four conditions for that outcome: AI must preserve meaningful regulatory decisions for learners; teachers must mediate the allocation of those decisions; support must be calibrated and faded as learners' regulatory readiness changes; and school resources must sustain the mediation routines over time. These conditions explain why identical AI tools can yield different educational outcomes across classrooms.

The framework makes three practical and theoretical contributions. First, the Learner Regulatory Opportunity (LRO) index operationalizes decision ownership as an observable

construct. Second, the Frame-Constrain-Interrogate-Fade protocol translates broad calls for human-centered AI into concrete classroom moves. Third, the framework treats equity as a pathway linking resources, teacher capacity, mediation fidelity, and learner opportunity. Together, these contributions identify two central tensions: personalization versus autonomy and efficiency versus productive effort. They also reveal a less visible equity problem: reliable AI access can coexist with limited opportunities to practice independent regulation.

The implications are therefore specific. For policymakers and school leaders, AI adoption should be accompanied by teacher professional development focused on mediation and verification, continuity of access rather than device counts alone, and assessment of independent transfer. For teachers, the practical routine is straightforward: require an initial attempt, limit substantive help, treat AI output as contestable, and schedule no-AI transfer. For researchers, the priorities are to validate LRO, test FCIF against AI-as-usual, examine readiness-dependent support intensity, and investigate whether resource conditions operate through mediation capacity. Ultimately, every AI integration decision can be tested with one question: What can the student now plan, monitor, judge, and complete without the system that they could not before? If assisted performance rises while independent regulatory competence does not, the technology may be personalizing without educating.

Classrooms would be assigned to two conditions: AI-as-usual and T-MAIR-mediated AI. Both conditions receive equal access to the same AI functions. In the T-MAIR condition, teachers use Frame-Constrain-Interrogate-Fade routines, including a learner-stated goal, an independent attempt before substantive assistance, an explicit verification move for major AI outputs, and scheduled no-AI transfer tasks. The comparison condition uses the platform under ordinary classroom practice without the mandated mediation sequence. Cluster allocation at the classroom or school level is preferable to student-level randomization because regulatory routines are visible and contamination within a class would be difficult to prevent. Fidelity should be observed rather than inferred from assignment.

Table 4. Research priorities for testing T-MAIR

Priority	Study	Primary question	Indicative horizon
High	LRO validation	Does LRO show reliability, construct validity, and predictive validity for transfer?	1–2 years
High	Cluster trial	Does T-MAIR improve SRL and delayed no-AI transfer versus AI-as-usual?	2–3 years
High	Readiness × fading experiment	Does optimal support intensity vary with baseline readiness and task difficulty?	1–2 years
Medium	Cross-cultural validation	Do mediation, fading, and transfer operate similarly across educational contexts?	3–4 years
Medium	Teacher-development trial	Can mediation-focused professional development improve fidelity and student outcomes?	2–3 years

A three-level longitudinal model is appropriate because repeated observations are nested within students, students within classrooms, and classrooms within schools. The confirmatory model should estimate post-intervention and delayed-transfer outcomes as functions of baseline score, condition, grade band/readiness, and resource context, including condition $\times$ readiness and condition $\times$ resource interactions. Mediation analysis can then test whether condition effects operate through LRO, verification behavior, or help-seeking latency. With a small number of schools, school-level random effects and cross-level estimates should be interpreted cautiously; a larger subsequent trial would be needed to obtain stable institution-level variance estimates.

The five propositions are restated here in parallel form for ease of testing:

Regulatory preservation means keeping target planning, monitoring, evaluation, and adaptation decisions available for learner practice. Regulatory displacement occurs when AI makes those decisions before the learner has an opportunity to do so. Productive metacognitive friction is the difficulty that prompts monitoring and reflection without overwhelming the learner. Developmental fading is the contingent reduction of external support as regulatory readiness increases. Independent transfer is the successful regulation of comparable work when supporting AI is unavailable.

P1. Higher Learner Regulatory Opportunity (LRO) will predict greater growth in metacognitive regulation and delayed no-AI transfer.

P2. Higher-quality teacher mediation will change the effect of identical AI support by strengthening learner-generated planning, monitoring, justification, and reflection.

P3. The relationship between scaffolding intensity and metacognitive growth will depend on regulatory readiness and task difficulty, producing an optimal range rather than a universal main effect.

CONCLUSIONS

The central question for K-12 AI integration is not whether artificial intelligence can personalize assistance, but whether that assistance leaves learners more capable of regulating comparable work when assistance is reduced or removed. The literature synthesized in this paper does not support either technological optimism or blanket rejection. AI-supported self-regulated learning can yield meaningful benefits, yet the pattern is heterogeneous across regulatory phases, intervention designs, developmental levels, and outcome measures. This heterogeneity is not merely a problem for implementation; it is a theoretical clue. It suggests that the educational effect of AI depends, in part, on who retains responsibility for the decisions that constitute learning.

T-MAIR offers a mechanism-based account of that responsibility. The framework positions teacher mediation between AI affordances and learner outcomes and specifies four observable moves: Frame establishes an independent learning target and success criterion; Constrain protects opportunities for learner-generated work before providing substantive assistance; Interrogate turns AI output into an object for verification, explanation, comparison, or critique; and Fade progressively returns regulatory responsibility to the learner through reduced prompting and no-AI transfer. The Learner Regulatory Opportunity (LRO) construct makes this allocation of responsibility observable, while developmental readiness and school resource continuity explain why the same mediation routine may not operate identically across learners or contexts.

The framework also advances an equity argument. Equitable AI integration should not be defined solely by whether students have access to comparable tools. Students should also receive comparable opportunities to practice planning, monitoring, judgment, help-seeking, and reflection. A well-resourced classroom can still foster regulatory dependence if AI routinely makes these decisions, whereas a resource-constrained classroom may protect meaningful learner agency through carefully timed, teacher-mediated use. Equity therefore concerns the distribution of high-

quality regulatory opportunities as well as material access.

T-MAIR consequently shifts the outcome criterion from assisted performance to durable competence. Immediate accuracy, completion, engagement, or fluent AI interaction can all be educationally useful, but none alone establishes independent regulation. A stronger test asks whether learners increasingly generate and justify the pivotal decisions themselves and whether those processes transfer to structurally comparable tasks without AI. The five propositions advanced in this paper translate that claim into a research program: LRO should predict transfer; teacher mediation should alter the function of identical AI support; optimal scaffolding should depend on readiness and task difficulty; resource conditions should influence outcomes, in part, through mediation capacity; and similar assisted performance should not be assumed to represent similar learning.

The contribution of T-MAIR is deliberately testable rather than technologically prescriptive. It does not claim that teacher mediation is always beneficial, that independence is always the preferred endpoint, or that AI should be minimized categorically. Instead, it proposes a measurable principle: AI is educationally successful when its support helps learners assume greater regulatory responsibility over time. The proposed LRO index, FCIF protocol, fidelity measures, and multilevel empirical design provide concrete routes for testing whether that principle holds across subjects, grade bands, and resource contexts.

LIMITATION & FURTHER RESEARCH

This paper has several limitations that constrain the strength of its claims. First, it is a structured conceptual synthesis rather than a systematic review or meta-analysis. The empirical literature served as mechanism-revealing anchor evidence and counterexamples rather than as a complete census of K–12 AI research or as a basis for a pooled causal estimate. Accordingly, the propositions should be understood as theoretically grounded hypotheses, not established causal relationships. Future reviews should preregister search and screening procedures, document exclusions, use independent reviewers, assess study quality, and examine whether conclusions are robust across developmental levels and research designs.

Second, the evidence base remains uneven across geographic, institutional, and technological contexts. Much school-based research is conducted in settings with relatively stable digital infrastructure, while low-connectivity, rural, multilingual, and otherwise resource-constrained contexts may organize AI use, teacher mediation, and shared-device practices differently. T-MAIR therefore requires cross-cultural and cross-resource validation rather than assuming that a single model of learner autonomy applies universally. In some educational settings, teacher authority and co-regulation may remain valued outcomes rather than merely transitional supports. Future studies should examine which forms of regulatory responsibility are transferable across contexts and which are culturally or institutionally specific.

Third, developmental readiness cannot be reliably inferred from grade level alone. Learners of the same age may differ substantially in prior knowledge, executive control, motivation, language proficiency, AI literacy, disability-related needs, and opportunities for support outside school. The proposed LRO index is also likely to require subject-sensitive coding because a meaningful regulatory decision point in mathematics may differ from those in writing, science inquiry, or collaborative problem solving. Before LRO is used as a stable outcome or mediator, research should establish inter-rater reliability, construct validity, sensitivity to task difficulty, and predictive validity for delayed no-AI transfer.

Fourth, implementation creates practical and ethical constraints. High-fidelity teacher mediation can increase planning and monitoring demands, while continuous inspection of student–AI interactions may create privacy and governance risks. The framework should therefore be tested

using minimum-effective mediation rather than assuming that more observation is better. Sampling interaction episodes, using privacy-preserving process traces, applying low-burden fidelity rubrics, and collecting targeted transfer tasks may provide a more scalable balance between research precision and classroom feasibility. Future implementation studies should also examine teacher workload, professional learning requirements, platform configurability, and the sustainability of mediation routines over an academic year, not just during short interventions.

The highest-priority empirical program follows directly from these limitations. A cluster-randomized trial should compare T-MAIR-mediated AI with AI-as-usual while holding core AI functions constant. A factorial study should manipulate scaffolding intensity and baseline regulatory readiness to test the predicted readiness-dependent optimum. Longitudinal studies should determine whether changes in LRO precede gains in metacognitive regulation and delayed no-AI transfer. Mixed-methods work should combine process traces, observations, student think-alouds, and teacher interviews so that statistical differences can be linked to the mechanisms that produced them. Disconfirming cases should be actively examined: strong transfer in low-resource settings or weak transfer despite high mediation fidelity would be theoretically valuable because such cases could reveal omitted mechanisms and refine T-MAIR. Together, these studies can determine whether the framework functions not simply as a persuasive conceptual model but as a reliable and generalizable account of how AI can be used to develop learner autonomy rather than dependence.

REFERENCES

- Agustin, M. C. (2026). AI as a co-regulator: Relational design for strengthening self-regulated learning. *Frontiers in Education*, 11, 1761602. <https://doi.org/10.3389/feduc.2026.1761602>
- Azevedo, R., Chauncey, A., Burkett, C., & Fike, A. (2009). MetaTutor: A metacognitive tool for enhancing self-regulated learning. In *Cognitive and metacognitive educational systems—Papers from the AAAI Fall Symposium* (pp. 14–19). AAAI. <https://doi.org/10.3233/978-1-60750-028-5-635>
- Azevedo, R., & Hadwin, A. F. (2005). Scaffolding self-regulated learning and metacognition: Implications for the design of computer-based scaffolds. *Instructional Science*, 33, 367–379. <https://doi.org/10.1007/s11251-005-1272-9>
- Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185–205). MIT Press. <https://doi.org/10.7551/mitpress/4561.003.0011>
- Boulhrir, T., Ghreir, H., Hamash, M., & Robert, M. (2026). Artificial intelligence in education: Mapping adaptive learning and learning analytics in K-12 online, virtual, and distance learning. *The International Review of Research in Open and Distributed Learning*, 27(2), 122–148. <https://doi.org/10.19173/irrodl.v27i2.9370>
- Chang, W.-L., & Sun, J. C.-Y. (2024). Evaluating AI's impact on self-regulated language learning: A systematic review. *System*, 126, 103484. <https://doi.org/10.1016/j.system.2024.103484>
- Chiu, T. K. F. (2024). A classification tool to foster self-regulated learning with generative artificial intelligence by applying self-determination theory: A case of ChatGPT. *Educational Technology Research and Development*, 72, 2401–2416. <https://doi.org/10.1007/s11423-024-10366-w>
- Dejacto, L. B., Valmorida, F. M. S., & Naparan, G. B. (2023). Navigating nuances: Intermediate teachers' experiences handling Generation Z learners. *Journal of Elementary and Secondary School*, 1(2), 40–53. <https://doi.org/10.31098/jess.v1i2.1855>
- Dent, A. L., & Koenka, A. C. (2016). The relation between self-regulated learning and academic achievement across childhood and adolescence: A meta-analysis. *Educational Psychology Review*, 28, 425–474. <https://doi.org/10.1007/s10648-015-9320-8>
- Dignath, C., & Büttner, G. (2008). Components of fostering self-regulated learning among students: A meta-analysis on intervention studies at primary and secondary school level. *Metacognition and Learning*, 3, 231–264. <https://doi.org/10.1007/s11409-008-9029-x>

- Donker, A. S., de Boer, H., Kostons, D., Dignath van Ewijk, C. C., & van der Werf, M. P. C. (2014). Effectiveness of learning strategy instruction on academic performance: A meta-analysis. *Educational Research Review*, 11, 1–26. <https://doi.org/10.1016/j.edurev.2013.11.002>
- 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. <https://doi.org/10.1177/1529100612453266>
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34(10), 906–911. <https://doi.org/10.1037/0003-066X.34.10.906>
- Fütterer, T., Bardach, L., Kuhn, J., Keller, S. D., & Gerjets, P. (2026). Enhancing school students' self-regulated learning through generative AI support: A randomized controlled trial. *Educational Psychology Review*, 38, 42. <https://doi.org/10.1007/s10648-026-10133-8>
- Greene, J. A., & Azevedo, R. (2007). A theoretical review of Winne and Hadwin's model of self-regulated learning: New perspectives and directions. *Review of Educational Research*, 77(3), 334–372. <https://doi.org/10.3102/003465430303953>
- Guan, R., Raković, M., Chen, G., & Gašević, D. (2025). How educational chatbots support self-regulated learning? A systematic review of the literature. *Education and Information Technologies*, 30, 4493–4518. <https://doi.org/10.1007/s10639-024-12881-y>
- Hadwin, A. F., Järvelä, S., & Miller, M. (2011). Self-regulation, co-regulation, and socially shared regulation of learning. *Educational Psychologist*, 46(2), 65–84. <https://doi.org/10.1080/00461520.2011.538686>
- Hargittai, E. (2010). Digital na(t)ives? Variation in Internet skills and uses among members of the "net generation." *Sociological Inquiry*, 80(1), 92–113. <https://doi.org/10.1111/j.1475-682X.2009.00317.x>
- Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of Educational Research*, 77(1), 81–112. <https://doi.org/10.3102/003465430298487>
- Horvers, A., Kooi, R., Knoop-van Campen, C. A. N., Dijkstra, R., Baars, M., & Molenaar, I. (2024). How does co-regulation with adaptive learning technologies affect primary school students' goal-setting, regulation of practice behavior and learning outcomes? *Frontiers in Education*, 9, 1435483. <https://doi.org/10.3389/feduc.2024.1435483>
- Iterbeke, K., De Witte, K., & Schelfhout, W. (2021). The effects of computer-assisted adaptive instruction and elaborated feedback on learning outcomes: A randomized control trial. *Computers in Human Behavior*, 120, 106666. <https://doi.org/10.1016/j.chb.2020.106666>
- Järvelä, S., & Hadwin, A. F. (2013). New frontiers: Regulating learning in CSEL. *Educational Psychologist*, 48(1), 25–39. <https://doi.org/10.1080/00461520.2012.748006>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist*, 41(2), 75–86. https://doi.org/10.1207/S15326985EP4102_1
- Klar, M. (2025). Using ChatGPT is easy, using it effectively is tough? A mixed methods study on K–12 students' perceptions, interaction patterns, and support for learning with generative AI chatbots. *Smart Learning Environments*, 12, 32. <https://doi.org/10.1186/s40561-025-00385-2>
- Kong, S.-C., & Yang, Y. (2024). A human-centered learning and teaching framework using generative artificial intelligence for self-regulated learning development through domain knowledge learning in K–12 settings. *IEEE Transactions on Learning Technologies*, 17, 1588–1599. <https://doi.org/10.1109/TLT.2024.3392830>
- Li, W., Xu, Y., Yao, L., & Liu, Y. (2026). Exploring the influence of AI self-study rooms on K–12 learners' motivation, self-regulation, enjoyment, and engagement. *Frontiers in Psychology*, 17,

1768389. <https://doi.org/10.3389/fpsyg.2026.1768389>
- Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences*, 13(4), 410. <https://doi.org/10.3390/educsci13040410>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- Manaig, K. A., Yazon, A. D., Buama, C. A. C., & Bonganciso, R. T. (2025). Teachers' awareness in artificial intelligence and digital competence in the workplace. *Advanced Journal of STEM Education*, 3(1), 27–44. <https://doi.org/10.31098/ajosed.v3i1.3302>
- Molenaar, I. (2022). Towards hybrid human-AI learning technologies. *European Journal of Education*, 57(4), 632–645. <https://doi.org/10.1111/ejed.12527>
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 125–173). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)60053-5](https://doi.org/10.1016/S0079-7421(08)60053-5)
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). AI literacy: Definition, teaching, evaluation and ethical issues. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
- Ng, D. T. K., Tan, C. W., & Leung, J. K. L. (2024). Empowering student self-regulated learning and science education through ChatGPT: A pioneering pilot study. *British Journal of Educational Technology*, 55(4), 1328–1353. <https://doi.org/10.1111/bjet.13454>
- Nguyen, H., Chen, H.-L., & Asfani, K. (2026). Supporting self-regulated learning in the age of artificial intelligence: A comparative systematic review of earlier AI and generative AI. *Journal of Educational Computing Research*. Advance online publication. <https://doi.org/10.1177/07356331261450647>
- Orhani, S. (2023). Robots assist or replace teachers in the classroom. *Journal of Elementary and Secondary School*, 1(1), 29–41. <https://doi.org/10.31098/jess.v1i1.1418>
- Panadero, E. (2017). A review of self-regulated learning: Six models and four directions for research. *Frontiers in Psychology*, 8, 422. <https://doi.org/10.3389/fpsyg.2017.00422>
- Pintrich, P. R. (2000). The role of goal orientation in self-regulated learning. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation: Theory, research, and applications* (pp. 451–502). Academic Press. <https://doi.org/10.1016/B978-012109890-2/50043-3>
- Pintrich, P. R., & De Groot, E. V. (1990). Motivational and self-regulated learning components of classroom academic performance. *Journal of Educational Psychology*, 82(1), 33–40. <https://doi.org/10.1037/0022-0663.82.1.33>
- Salah, M., Al-Busaidi, R., Al-Ghawi, A., & Shawky, W. (2026). The regulation paradox: Agentic AI, bounded autonomy, and self-regulated learning in higher education. *Frontiers in Education*, 11, 1902833. <https://doi.org/10.3389/educ.2026.1902833>
- Schraw, G., & Dennison, R. S. (1994). Assessing metacognitive awareness. *Contemporary Educational Psychology*, 19(4), 460–475. <https://doi.org/10.1006/ceps.1994.1033>
- Theobald, M. (2021). Self-regulated learning training programs enhance university students' academic performance, self-regulated learning strategies, and motivation: A meta-analysis. *Contemporary Educational Psychology*, 66, 101976. <https://doi.org/10.1016/j.cedpsych.2021.101976>
- Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, 10, 15. <https://doi.org/10.1186/s40561-023-00237-x>
- Valenzona, R. B. (2026). Instructional materials enhancing high school learners' engagement in research skills. *Journal of Elementary and Secondary School*, 4(1), 18–35. <https://doi.org/10.31098/jess.v4i1.3660>
- Veenman, M. V. J., Van Hout-Wolters, B. H. A. M., & Afflerbach, P. (2006). Metacognition and learning: Conceptual and methodological considerations. *Metacognition and Learning*, 1, 3–14. <https://doi.org/10.1007/s11409-006-6893-0>
- Wang, X., Zhang, L. J., & Zhang, Y. (2026). Generative artificial intelligence-supported self-regulated learning (GenAI-SRL) in L2 writing: Scale development, validation, and short-form construction. *Metacognition and Learning*, 21, 33. <https://doi.org/10.1007/s11409-026->

09481-1

- Warschauer, M., & Matuchniak, T. (2010). New technology and digital worlds: Analyzing evidence of equity in access, use, and outcomes. *Review of Research in Education*, 34(1), 179–225. <https://doi.org/10.3102/0091732X09349791>
- Winne, P. H., & Hadwin, A. F. (1998). Studying as self-regulated learning. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 277–304). Lawrence Erlbaum Associates. <https://doi.org/10.4324/9781410602350-19>
- 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. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>
- Xia, Q., Liu, Q., Tlili, A., & Chiu, T. K. F. (2026). A systematic literature review on designing self-regulated learning using generative artificial intelligence and its future research directions. *Computers & Education*, 240, 105465. <https://doi.org/10.1016/j.compedu.2025.105465>
- Xu, J., Luo, Y., Wang, C., Wang, M., & Wu, Y. (2026). AI support in self-regulated learning: A decade of technological evolution and meta-analysis. *British Journal of Educational Technology*, 57(5), 1191–1220. <https://doi.org/10.1111/bjet.70058>
- Xu, X., Qiao, L., Cheng, N., Liu, H., & Zhao, W. (2025). Enhancing self-regulated learning and learning experience in generative AI environments: The critical role of metacognitive support. *British Journal of Educational Technology*, 56, 1842–1863. <https://doi.org/10.1111/bjet.13599>
- Zhang, J., & Derakhshan, A. (2025). Integrating GPT-4o as an AI tool into K–12 contexts: Assessing its long-term effects on students' self-regulated learning and task engagement through latent growth curve modelling. *European Journal of Education*, 60(4), e70329. <https://doi.org/10.1111/ejed.70329>
- Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, 41(2), 64–70. https://doi.org/10.1207/S15430421TIP4102_2

APPENDIX

1. Minimum design: 12–16 classrooms, at least four schools, and both elementary and secondary cohorts. Use a pretest, repeated process measures, a post-test, and delayed no-AI transfer.
2. Student variables: age/grade, prior achievement, baseline SRL/metacognition, AI literacy, perceived utility, self-efficacy, and home technology access.
3. Teacher variables: AI literacy, years of experience, perceived workload, mediation fidelity, and professional-development exposure.
4. Platform/process variables: time-on-task, first-help latency, number and function of prompts, independent attempts, verification events, revisions after feedback, independent retries, use of adaptation features, and the proposed LRO index. Preserve event sequence and timestamps so that regulatory episodes can be reconstructed rather than reduced to totals.
5. School variables: device availability, bandwidth reliability, class size, technical support, timetable flexibility and AI policy.
6. Primary outcomes: metacognitive regulation, independent transfer and learner autonomy. Secondary outcomes: achievement, engagement, motivation and enjoyment.
7. Recommended analysis: multilevel pretest-adjusted models with condition × grade band and condition × resource-context interactions; mediation analysis through help-seeking and verification behavior; qualitative joint display for mechanism interpretation.

Table 5. Teacher-mediation fidelity rubric for a future T-MAIR field test

Dimension	0 = Absent	1 = Partial	2 = Clear implementation
Frame	No explicit	Goal stated by	Student articulates

	independent goal	teacher only	goal and success criterion
Constrain	AI used before student attempt	Attempt requirement inconsistent	Independent attempt precedes substantive AI help
Interrogate	Output accepted without checking	General checking prompt	Evidence, explanation or comparison explicitly required
Fade	No reduction of support	Occasional no-AI task	Planned transfer plus responsive removal of prompts

Table 6. Minimum measurement specification for a future T-MAIR field test

Level	Variable	Operational example	Role in analysis
Student	Baseline SRL	Validated age-appropriate scale	Moderator/covariate
Student	Independent transfer	Delayed no-AI assessment	Primary/secondary outcome
Process	First-help latency	Seconds from task start to first AI request	Behavioral mediator
Process	Verification behavior	Checks against evidence/source or self-explanation	Behavioral mediator
Teacher	Mediation fidelity	0–8 observation score	Implementation predictor
Teacher	AI literacy	Validated literacy/confidence measure	Moderator
Classroom	Class size	Enrolled students per class	Context covariate
School	Resource index	Devices, connectivity, support, training	Cross-level moderator
Process	Learner Regulatory Opportunity (LRO)	Learner-generated/justified regulatory decisions ÷ coded decision points	Proposed mediator / process indicator