VanLehn's Two Loops Aren't Enough: Toward Meta-Adaptive Control of Instructional Form

Bruce M. McLaren(✉)

Carnegie Mellon University, Pittsburgh, USA
bmclaren@cs.cmu.edu

Abstract. For nearly two decades, research on Intelligent Tutoring Systems (ITSs) has been guided by VanLehn's two-loop model, framing adaptation as optimizing what learners do next through task selection and step-level feedback. However, this framework assumes that instructional form – learning technology, epistemic activity, modality, and assessment stance – is fixed. This assumption no longer holds in contemporary learning ecologies, where learners move fluidly across diverse learning environments, including tutors, games, simulations, virtual and augmented reality, and conversational agents. I propose a third, meta-adaptive loop that elevates instructional form to a first-class adaptive variable, shifting adaptation from selecting tasks within activities to selecting how learners engage at each moment. Drawing on learning theory (e.g., ICAP, productive failure, desirable difficulties), I argue that *how learners engage* often matters as much as *what they do*. Recent advances in LLM-based AI make principled cross-form orchestration increasingly feasible, enabling coordination across heterogeneous systems. This reframes AIED from optimizing performance within activities to optimizing the structure of learning itself, with implications for future architectures, benchmarks, and study designs.

Keywords: Intelligent Learning Technology · Meta-Adaptation · Instructional Form · Instructional Orchestration · Learning Ecologies · Inner Loop · Outer Loop

1 Introduction

For nearly two decades, the dominant architectural metaphor for Intelligent Tutoring Systems (ITSs) has been VanLehn's two-loop model, distinguishing between an outer loop that selects tasks and an inner loop that provides step-level feedback and guidance during problem solving (VanLehn, 2006, 2011). This framework has strongly shaped ITS research and development. Its impact extends beyond system design to how the field conceptualizes adaptation itself (cf. Aleven et al., 2016).

Implicit in this model, however, is a key assumption: that instructional form – defined here as the structure of learning activity (e.g., problem solving, example study, simulation, collaboration, or assessment), rather than the specific content of tasks – is fixed. Adaptation occurs *within* a form, but not *across* forms. This paper argues that instructional form is a runtime control variable. In contemporary learning environments, the

E. G. Blanchard et al. (Eds.): AIED 2026, CCIS 3032, pp. 355–367, 2026.
https://doi.org/10.1007/978-3-032-29791-4_28

adaptive question is often not just "which problem next?" but "how should this learner engage right now – and why?" I argue that the dominant paradigm of AIED – optimizing what learners do next – is no longer sufficient. Instead, intelligent learning systems must also optimize how learners engage, by treating instructional form itself as a primary object of adaptive control. This shift reframes adaptation from selecting tasks within a fixed activity to selecting the activity itself, based on learner state, context, and pedagogical goals.

Since the introduction of the two-loop model, AIED has evolved from predominantly single-system tutors to heterogeneous learning ecologies encompassing a wide range of learning technologies (Aleven et al., 2023). These include digital learning games (Mayer, 2019; McLaren & Nguyen, 2023), interactive simulations and microworlds (de Jong & van Joolingen, 1998; Yaron et al., 2010), worked and erroneous examples (McLaren et al., 2016; Van Gog et al., 2011), embodied and tangible learning environments (Arroyo et al., 2022; Lindgren & Johnson-Glenberg, 2013; Nathan, 2021), immersive virtual and augmented reality environments (Makransky & Petersen, 2019; Yang, 2023), conversational agents and pedagogical chatbots (Alrajhi, 2024; Yusuf et al., 2025), and adaptive and stealth assessment systems (Shute & Ventura, 2013). This expansion also includes collaborative and socially interactive learning systems, adaptive peer tutoring, and analytics-supported scripted collaboration, which introduce qualitatively different epistemic activities and adaptive demands than individual learning systems (Dillenbourg, 1999; Stahl et al., 2006). Learners increasingly move across these instructional forms within and across learning sessions, often without principled, learner-sensitive orchestration.

Yet VanLehn's two-loop model offers no representational or theoretical account of how such between-system instructional decisions should be made. I therefore propose a third, meta-adaptive loop – operating above the inner and outer loops – that treats instructional form as an explicit runtime control variable and enables principled orchestration across learning technologies.

2 The Two-Loop Model and its Intended Scope

VanLehn's two-loop account characterizes adaptivity as step-level regulation during problem solving (inner/step loop) coupled with task selection based on a learner model (outer/task loop) (VanLehn, 2006, 2011). His later regulative taxonomy clarifies that these loops differ mainly by granularity and timescale, not by instructional paradigm (VanLehn, 2016). Critically, the paradigm itself is treated as fixed: the loops operate **within** a chosen form (e.g., a problem-solving tutor, a learning game) rather than selecting among forms (e.g., tutor vs. simulation vs. game vs. collaboration vs. assessment). The meta-adaptive loop proposed here extends the regulative perspective by making instructional form – how the learner is asked to learn – an explicit runtime decision. Related ideas have appeared in the ITS literature. For example, Pavlik et al. (2013) noted that student models may support pedagogical decisions at multiple levels of granularity, from individual problems to broader curriculum units and domains, but did not develop an explicit architecture for adapting instructional form itself.

3 Learning Theory Motivates Instructional Form as an Adaptive Variable

Decisions about instructional form – such as whether learners should observe worked examples, attempt problem solving, explore simulations, engage in game-based practice, or interact through dialogue – have long been central to learning theory. Across theoretical traditions, these decisions shape learning outcomes, engagement, and transfer.

One influential account is the ICAP framework, which distinguishes among **I**nteractive, **C**onstructive, **A**ctive, and **P**assive modes of engagement and predicts systematic differences in learning outcomes across them (Chi & Wylie, 2014). Importantly, ICAP does not prescribe a single optimal mode of instruction. Rather, it emphasizes that the effectiveness of an activity depends on how learners engage with it, as well as on learner characteristics and instructional context. From this perspective, decisions about whether learners should generate explanations, solve problems, or engage in interactive dialogue are theoretically consequential – and potentially adaptive.

Multiple complementary lines of theory converge on the same point: *instructional effectiveness depends on matching forms of engagement to learner state and timing*. Research on productive failure (Kapur, 2008) and guided discovery (de Jong & van Joolingen, 1998; Kirschner et al., 2006) shows that exploration benefits learners with sufficient prior knowledge, constrained tasks, and consolidation, but can overload novices when these conditions are absent. A meta-adaptive loop operationalizes these boundary conditions by shifting learners into exploratory simulations only when indicators of readiness are met and reverting to guided instruction when exploration becomes unproductive. Complementary work on instructional sequencing further shows that learning outcomes depend critically on when learners engage in examples, practice, or instruction, rather than on any single activity in isolation (Koedinger et al., 2013; Renkl, 2014). Similarly, the *Assistance Dilemma* highlights the importance of timing and learner state in deciding when to provide guidance and when to withhold it (Koedinger & Aleven, 2007).

Other theoretical traditions point to additional dimensions of instructional form. Work on embodied cognition argues that physical interaction and sensorimotor engagement can ground conceptual understanding in ways that differ fundamentally from abstract symbolic manipulation (Lindgren & Johnson-Glenberg, 2013; Nathan, 2021). Research on assessment and motivation shows that how learners are evaluated – through explicit testing or stealth assessment embedded in activity – can influence anxiety, persistence, and engagement, independently of measurement accuracy (Shute & Ventura, 2013; Pekrun, 2006).

Despite this theoretical richness, most intelligent learning systems treat instructional form as a fixed design choice. Existing architectures support adaptation of tasks and feedback but provide no explicit mechanism for reasoning about epistemic activity, modality, or assessment stance. As a result, theoretically consequential instructional decisions remain largely outside adaptive control. Together, these perspectives suggest that instructional form is a dynamically adjustable dimension of learning.

4 Limits of the Two-Loop Model in Contemporary Learning Ecologies

As learning technologies diversify, the two-loop model remains effective for within-form adaptation but leaves between-form decisions unmodeled in increasingly heterogeneous learning ecologies (Dede, 2014).

First, the two-loop model cannot account for adaptation across multiple instructional forms that target the same competencies and skills but differ in modality, epistemic activity, and learner experience. Decisions such as whether a learner should engage with a digital learning game, an intelligent tutor, or an exploratory simulation are pedagogy-level choices, not task-level ones. Collaborative learning further illustrates the gap: choosing whether learning is individual or collaborative (and with whom/under what script) changes the unit of adaptation and cannot be reduced to task or step selection.

Second, the model offers no mechanism for dynamically shifting between instructional and assessment modes. The growing use of stealth assessment embedded in games or simulations challenges the traditional separation between instruction and evaluation (Shute & Ventura, 2013). While outer-loop mechanisms can select assessment items, they cannot represent the assessment stance itself as an adaptive variable. Empirical work on orchestration load suggests that managing such decisions places substantial cognitive and logistical demands on instructors (Dillenbourg, 2013), further motivating the need for computational support at this level of control.

Third, the model cannot represent adaptation among pedagogical mechanisms such as worked examples, erroneous examples, tutored problem solving, and exploratory learning. A substantial body of empirical work shows that these approaches differ meaningfully in efficiency, effectiveness, and learner experience depending on prior knowledge and context (de Jong & van Joolingen, 1998; McLaren et al., 2016; Van Gog et al., 2011). Yet within the two-loop framework, the choice among these mechanisms is typically fixed at design time, with adaptation limited to selecting the next problem or hint.

Prior work has explored adaptive selection among learning activities such as worked examples and problem solving (e.g., Shareghi Najar et al., 2016; Chen et al., 2020). However, these approaches operate within a single system, treating activities as task variants. In contrast, the meta-adaptive loop elevates instructional form itself – across tutors, simulations, games, collaboration, and assessment – to an explicit architectural control variable.

Fourth, a related and widely recognized challenge is the difficulty of sharing and operationalizing student models across learning systems. In the BLAP "Transferability" challenge, Baker (2019) identified the difficulty in transferring learner models between systems as a fundamental barrier to cumulative adaptation in AIED. Without mechanisms for cross-system learner state integration, each learning environment operates as an "adaptive island." The meta-adaptive loop proposed here provides an architectural locus for integrating such cross-form learner models and coordinating transitions among systems based on unified state representations. This challenge underscores that adaptation across systems requires not just interoperability, but architectural representation of cross-form control.

One might argue that a sufficiently rich outer-loop policy could encode form switching implicitly. However, doing so collapses qualitatively different instructional forms into a single task-selection mechanism, obscuring pedagogical intent and making instructional form an unmodeled side effect rather than an explicit target of adaptation.

As a result, contemporary learning systems risk optimizing within instructional forms that are poorly matched to learners' needs. A system can look highly adaptive – fine-grained hinting, precise sequencing – while still failing at the most consequential decision: *whether the learner is in the right instructional form at all.*

5 A Meta-Adaptive Loop: Concept and Architecture

The meta-adaptive loop is proposed as an explicit architectural layer that regulates instructional form – i.e., how learning occurs – distinct from both task selection and step-level guidance. Whereas the inner loop governs within-task interaction – providing feedback, hints, and guidance as learners act – and the outer loop governs which task or problem is selected next within a fixed instructional form, the meta-adaptive loop governs *how learning proceeds* by selecting the instructional form itself. Instructional form here refers to the learning activity format or pedagogical mechanism that structures engagement, such as worked examples versus problem solving, game-based practice versus direct instruction, exploratory simulation versus guided tutoring, embodied versus abstract representations, or stealth versus explicit assessment.

Mainstream learning technology architectures have not formalized instructional form itself as a runtime control layer distinct from task and step regulation. Instead, between-form decisions are typically hard-coded, externally sequenced, or implicitly folded into task selection policies. As a result, qualitatively different pedagogical mechanisms are collapsed into a single "next task" abstraction, obscuring pedagogical intent and limiting interpretability.

This reframing shifts adaptation from optimizing performance within activities to selecting the activities themselves.

The meta-adaptive loop operates above the inner and outer loops and regulates transitions among instructional forms based on learner state, contextual constraints, and pedagogical goals (see Fig. 1). Architecturally, the meta-adaptive controller consumes a cross-form state vector (knowledge estimates, engagement/affect signals, contextual constraints, and recent instructional history) and outputs (i) a selected instructional form and (ii) a handoff specification (target skills, time budget, and transferable learner state) to the invoked environment's inner/outer loops.

Formally, the meta-adaptive loop can be modeled as a constrained control policy. The **state** comprises a cross-form learner model (knowledge estimates, engagement/affect proxies), contextual constraints (time, stakes, setting), and an explicit switching cost. The **action** selects an instructional form and parameterizes the handoff (target skills, time budget, transferable learner state). The **policy** is governed by constraints on allowable forms, validated triggers, and rationale logging to ensure interpretability and auditability.

Individual systems retain their internal adaptive control, while the meta-adaptive loop coordinates transitions among them. The three loops differ in object of control: steps, tasks, and instructional form.

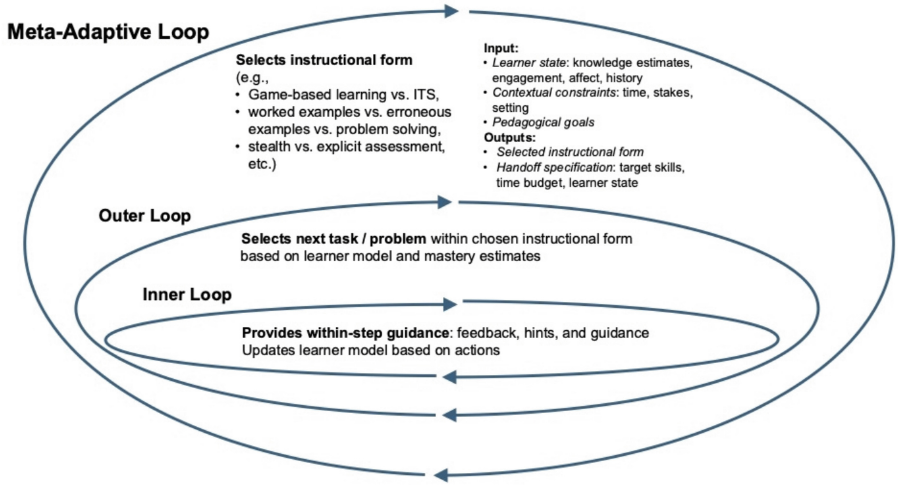

Fig. 1. The Meta-Adaptive Loop and its relationship with VanLehn's Outer and Inner Loops

The meta-adaptive loop may decide whether learners should observe worked or erroneous examples versus attempt problem solving; engage in game-based practice versus direct instruction; explore simulations versus receive guided tutoring; interact with embodied or immersive representations versus abstract symbolic ones; or be assessed implicitly through gameplay versus explicitly through formal testing.

Crucially, the meta-adaptive loop does not replace the inner and outer loops; individual learning technologies retain their internal adaptive mechanisms, while the meta-adaptive loop orchestrates transitions among them, enabling adaptation across instructional forms rather than only within them.

Table 1 summarizes the shift in what is treated as the object of adaptation.

Table 1. Contrasting Assumptions of the Two-Loop Model and the Meta-Adaptive Loop.

Two-Loop Model	Meta-Adaptive Loop Model
Operates *within* a system	Operates *across* systems
Assumes fixed pedagogy	Treats pedagogy (form) as adaptive
Adapts steps and tasks	Adapts instructional form
Tutor-centric	Ecosystem-centric

Table 2 illustrates that "next task" is an inadequate proxy for instructional form, which changes the interaction topology rather than task instance.

Instructional form differs from task selection in that it changes the topology of the interaction itself – whether the learner is observing, generating, collaborating, exploring, or being evaluated – rather than merely which instance of a task is presented. This distinction is structural, not parametric. Modeling form transitions as task selections

Table 2. Common choices between different pedagogies and instructional forms and the limitations of VanLehn's inner- and outer-loop architecture in selecting between them.

Learning technology / approach	Adaptive decision required	Why the two-loop model is insufficient
Intelligent Tutoring Systems (ITS)	**How to scaffold steps; which problem to assign**	Well captured by inner and outer loops
Digital Learning Games	**Whether to learn through gameplay vs. tutoring**	Assumes fixed instructional form; cannot represent game vs. tutor choice
Interactive simulations / microworlds	**Explore vs. guided problem solving**	Exploration vs. guided instruction is a pedagogy-level decision, not task selection
Worked examples	**Observe vs. practice**	Choice of pedagogy lies outside inner/outer loops
Erroneous examples	**Learn from errors vs. correct examples**	Pedagogical mechanism is fixed at design time in two-loop models
Tutored problem solving	**When to shift from examples to problem solving**	Transition among mechanisms not representable within inner/outer-loop control
Embodied / tangible learning environments	**Physical vs. symbolic interaction and learning**	Modality choice is only awkwardly expressible within a single tutor
Stealth assessment	**Assess implicitly vs. explicitly**	Assessment stance is external to two-loop adaptation
Formal testing	**Low-stakes vs. high-stakes evaluation**	Assessment treated as separate from instruction
Collaborative learning systems (CSCL, peer tutoring)	**Individual vs. collaborative activity; group composition; interaction structure**	Changes unit of learning and epistemic activity; not representable as task or step selection
Hybrid learning ecosystems	**Orchestration across multiple systems**	No mechanism for between-system control

collapses changes in interaction topology into content variation, obscuring pedagogical intent and limiting interpretability.

6 Theoretical Motivation and Illustrative Meta-Adaptive Decisions

The motivation for a meta-adaptive loop is fundamentally theoretical rather than technological. Across learning systems, instructional decisions about epistemic activity, modality, and assessment stance are known to shape learning outcomes – yet current architectures lack a principled way to operationalize these distinctions at runtime.

Importantly, instructional form need not be limited to canonical learning activities such as problem solving or example study. A meta-adaptive loop could also regulate transitions to non-traditional states, such as brief breaks, reflective pauses, or "time-for-telling" episodes (Schwartz & Bransford, 1998). It may even incorporate off-task activities that support motivation, consolidation, or cognitive reset.

6.1 Choosing Epistemic Activity and Instructional Form

The ICAP framework provides a clear motivation for meta-level adaptation (Chi & Wylie, 2014). Because different instructional forms elicit different modes of engagement, a meta-adaptive loop can select activities that promote deeper engagement when learners are ready, or more supported forms when they are not.

For example, a meta-adaptive system might initially assign worked examples to a novice learner exhibiting high error rates, transition to tutored problem solving as understanding stabilizes, and later select a digital learning game to promote fluency and sustained engagement. This sequence operationalizes ICAP distinctions while also aligning with empirical findings on example–problem fading and instructional efficiency (Van Gog et al., 2011; McLaren et al., 2016). While outer-loop mechanisms can sequence problems, they do not capture this shift in epistemic activity.

From a meta-adaptive perspective, collaboration itself can be treated as an instructional form that is selectively invoked. For example, a learner who repeatedly fails to explain reasoning in solo problem solving might be paired with a peer for structured collaborative explanation, while another learner showing off-task behavior in group work might be shifted back to individual tutoring. Such decisions align with theories of socially shared regulation and collaborative knowledge construction (Dillenbourg, 1999; Fischer et al., 2013; Roschelle & Teasley, 1995) yet cannot be expressed within traditional inner- or outer-loop control, which presuppose a fixed interaction topology.

6.2 Choosing Exploration Versus Guidance

Exploration tends to help when learners have enough prior knowledge to generate plausible hypotheses, tasks are constrained, and exploration is followed by consolidation; otherwise it can overload novices and depress learning (Kapur, 2008; Kirschner et al., 2006; de Jong & van Joolingen, 1998).

Importantly, optimizing instructional form should not be equated with maximizing short-term engagement or reducing difficulty. A meta-adaptive loop must also account for desirable difficulties (Bjork & Bjork, 2011) and productive struggle (Kapur, 2008), which may temporarily reduce engagement while improving long-term learning outcomes. In this sense, meta-adaptive control is not about selecting the most immediately engaging activity, but about selecting the form that best supports learning over appropriate timescales, even when that entails challenge or frustration.

6.3 Modality and Representation

Work on embodied cognition motivates adaptation in representational modality (Lindgren & Johnson-Glenberg, 2013; Nathan, 2021). For instance, a learner might begin with an embodied or tangible activity to ground conceptual understanding, then transition to symbolic problem solving once core concepts are established. Such modality shifts cannot be expressed within traditional inner- or outer-loop control yet are theoretically grounded and pedagogically consequential.

6.4 Selection of Assessment Based on Affect

Assessment theory provides another motivation. Research on stealth assessment shows that embedding evaluation within gameplay can reduce anxiety and support motivation, particularly for learners who experience test-related stress (Shute & Ventura, 2013). A meta-adaptive loop can therefore select between stealth and explicit assessment based on affective indicators, treating assessment stance itself as an adaptive instructional decision.

For example, consider a learner whose performance data indicate adequate conceptual mastery but whose affective signals – such as elevated frustration, disengagement, or avoidance behaviors – suggest rising anxiety. A meta-adaptive loop could respond by temporarily shifting the learner from an explicitly assessed tutoring environment to a game-based or simulation-based activity that embeds assessment stealthily within low-stakes interaction. Conversely, for a learner exhibiting high confidence but low metacognitive calibration, the system might shift from stealth assessment to explicit testing to support self-monitoring. These decisions are not about which task to present next, but about selecting an assessment *stance* that best supports learning and motivation at that moment.

6.5 Choosing Reflection Modality

Finally, theories of dialogic learning and reflection motivate adaptation in interaction structure. Some learners may benefit from multi-turn Socratic dialogue, while others may require only brief prompts. Selecting between these interaction forms reflects a change in epistemic support rather than task selection and aligns with ICAP as well as emerging work on pedagogical conversational agents (Yusuf et al., 2025).

Across these cases, learning theory converges on a common insight: many of the most consequential instructional decisions concern *how* learners engage. The meta-adaptive loop provides an architectural locus for enacting these distinctions dynamically.

7 Design Implications for Intelligent Learning Systems

Introducing a meta-adaptive loop enables intelligent learning systems to make principled, theory-informed decisions about how learners should engage – decisions that existing architectures either hard-code or leave to human instructors. Designers must move beyond asking what problem a learner should solve next to asking what kind of learning activity is most appropriate at a given time. Instructional form becomes a runtime adaptive variable rather than a static design choice. Because learners differ widely in prior knowledge, affective responses, and sociocultural experiences, adapting instructional form may also be critical for supporting equity and avoiding one-size-fits-all pedagogical assumptions (cf. Hyde et al., 2019).

Recent advances in Large Language Models (LLMs) suggest that what was previously a design-time decision – selection of instructional form – can now be treated as a runtime adaptive choice. LLMs provide a practical mechanism for mediating across heterogeneous learning environments, making cross-form orchestration newly feasible. Yet analyses of LLMs' educational use emphasize the importance of transparency, human-centered design, and pedagogical alignment (Kasneci et al., 2023). Research on human–AI complementarity further suggests that AI systems are most effective when they augment, rather than replace, human pedagogical judgment (Holstein & Aleven, 2021). From this perspective, LLMs are best viewed as mediators within a constraint-driven meta-adaptive loop, supporting principled orchestration rather than acting as autonomous tutors.

LLMs may support the meta-adaptive loop by (i) interpreting heterogeneous signals across tools (performance traces, interaction patterns, affect proxies, context) – a capability that is difficult to encode using hand-authored cross-system rules – and (ii) generating structured rationales for form switches – critical for transparency (Kasneci et al., 2023). However, LLM recommendations should be bounded by an explicit policy layer (allowable forms, validated triggers, switching costs, maximum switch frequency) so that the meta-loop remains interpretable and auditable rather than purely model-driven (Holstein & Aleven, 2021).

The introduction of a meta-adaptive loop raises important ethical considerations. Biases in state estimation (e.g., affect or engagement inference) may systematically influence how learners are routed across instructional forms. Ensuring fairness requires validating sensing mechanisms and auditing form-selection policies. Transparency is also critical: learners and instructors should be able to understand and, where appropriate, contest these decisions. Together, these concerns highlight the need for interpretable policies and rationale logging in meta-adaptive systems.

8 Challenges and Open Research Questions

Implementing a meta-adaptive loop raises substantial technical and methodological challenges. From an engineering perspective, such systems must integrate learner models across heterogeneous platforms, reconcile different representations of knowledge and engagement, and operate under real-time constraints without destabilizing lower-level adaptive processes. Designing control policies that balance responsiveness with instructional coherence - avoiding excessive or premature switching between forms - remains an open problem.

Evaluating the effectiveness of meta-adaptation poses equally significant challenges. Traditional intelligent learning system evaluations assume a fixed instructional form, making it difficult to isolate the causal impact of form-level transitions. Studying a third loop requires experimental designs in which transitions between instructional forms are the manipulated variable, alongside benchmarks that distinguish the benefits of dynamic orchestration from well-designed static sequences. Questions of transparency, interpretability, and learner trust become especially salient when systems adapt not just what learners do, but how they are asked to learn. A key empirical question is whether meta-adaptive policies outperform well-designed fixed sequences, requiring randomized controlled studies comparing (i) static instructional sequences, (ii) adaptive within-form systems, and (iii) meta-adaptive cross-form policies.

9 Conclusion

VanLehn's two-loop model remains foundational for understanding adaptation within intelligent learning technologies, providing a powerful account of how systems select tasks and guide learners during problem solving. However, as AIED increasingly operates within heterogeneous learning ecologies, the assumption of a fixed instructional form becomes untenable. The most consequential adaptive decisions are often not about which problem to present next, but about **which instructional form should govern learning at a given moment**.

This paper has proposed a third, meta-adaptive loop that elevates instructional form – epistemic activity, modality, and assessment stance – to an explicit object of adaptive control. In doing so, it reframes the locus of adaptation from optimizing performance within activities to selecting the activities themselves. This shift aligns closely with decades of learning theory, which consistently shows that engagement structure, timing, and modality critically shape learning outcomes, yet have remained largely outside the scope of adaptive architectures.

If realized, meta-adaptive systems would move AIED beyond isolated adaptive tools toward coordinated learning ecosystems, capable of orchestrating transitions among tutors, simulations, games, collaborative activities, and assessment modes in a principled, learner-sensitive way. This raises important technical, methodological, and ethical challenges, including cross-system learner modeling, evaluation of form-level adaptation, and the need for transparency and fairness in form selection policies.

More broadly, this work invites the field to reconsider what intelligent learning systems should adapt. The central claim is not simply that we need more adaptivity, but that we must adapt the right thing. Without adapting instructional form, AIED risks becoming increasingly precise about what learners do next, while remaining fundamentally blind to how they should learn.

Disclosure of Interests. The author has no competing interests to declare that are relevant to the content of this article.

References

Aleven, V., Mavrikis, M., McLaren, B.M., Nguyen, H.A., Olsen, J.K., Rummel, N.: Six instructional approaches supported in AIED systems. In: du Boulay, B., Mitrovic, A., Yacef, K. (eds.) Handbook of Artificial Intelligence in Education. Chapter 9 (2023). https://doi.org/10.4337/9781800375413.00019

Aleven, V., McLaughlin, E.A., Glenn, R.A., Koedinger, K.R.: Instruction based on adaptive learning technologies. In: Mayer, R.E., Alexander, P. (eds.) Handbook of Research on Learning and Instruction. Routledge (2016)

Alrajhi, A.S.: Artificial intelligence pedagogical chatbots as L2 conversational agents. Cogent Educ. **11**(1), 2327789 (2024). https://doi.org/10.1080/2331186X.2024.2327789

Arroyo, I., Closser, A.H., Castro, F., Smith, H., Ottmar, E., Micciolo, M.: The wearable learning platform: a computational thinking tool supporting game design and active play. In: Technology, Knowledge and Learning. Springer Nature (2022)

Baker, R.S.: Challenges for the future of educational data mining: the Baker learning analytics prizes. J. Educ. Data Min. **11**(1), 1–17 (2019). https://doi.org/10.5281/zenodo.3554745

Bjork, R.A., Bjork, E.L.: Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In: Psychology and the Real World (2011).

Chen, X., Mitrovic, A., Mathews, M.: Learning from worked examples, erroneous examples, and problem solving: toward adaptive selection of learning activities. IEEE Trans. Learn. Technol. **13**(1), 135–149 (2020). https://doi.org/10.1109/TLT.2019.2896080

Chi, M.T.H., Wylie, R.: The ICAP framework: linking cognitive engagement to active learning outcomes. Educ. Psychol. **49**(4), 219–243 (2014). https://doi.org/10.1080/00461520.2014.965823

Dede, C.: The role of digital technologies in deeper learning. Students at the Center: Deeper Learning Research Series. Boston, MA: Jobs for the Future (2014).

Dillenbourg, P.: Design for classroom orchestration. Comput. Educ. **69**(1), 485–492 (2013)

de Jong, T., van Joolingen, W.R.: Scientific discovery learning with computer simulations of conceptual domains. Rev. Educ. Res. **68**(2), 179–201 (1998). https://doi.org/10.2307/1170753

Dillenbourg, P.: What do you mean by collaborative learning? In: Dillenbourg, P. (ed.) Collaborative-Learning: Cognitive and Computational Approaches, pp. 1–19. Elsevier, Oxford (1999)

Fischer, F., Kollar, I., Stegmann, K., Wecker, C.: Toward a script theory of guidance in computer-supported collaborative learning. Ed. Psych. **48**(1), 56–66 (2013)

Holstein, K., Aleven, V.: Designing for human–AI complementarity in K-12 education. AI Mag. **42**(2), 3–14 (2021)

Hyde, J.S., Bigler, R.S., Joel, D., Tate, C.C., van Anders, S.M.: The future of sex and gender in psychology: five challenges to the gender binary. Am. Psychol. **74**(2), 171–193 (2019). https://doi.org/10.1037/amp0000307

Kapur, M.: Productive failure. Cogn. Instr. **26**(3), 379–424 (2008). https://doi.org/10.1080/07370000802212669

Kasneci, E., et al.: Chat GPT for good? On opportunities and challenges of large language models for education. Learn. Individ. Differ. **103**, 102274. (2023). https://doi.org/10.1016/j.lindif.2023.102274

Kirschner, P.A., Sweller, J., Clark, R.E.: Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. Educ. Psychol. **41**(2), 75–86 (2006). https://doi.org/10.1207/s15326985ep4102_1

Koedinger, K.R., Booth, J.L., Klahr, D.: Instructional complexity and the science to constrain it. Science. **342**(6161), 935–937 (2013). https://doi.org/10.1126/science.1238056

Koedinger, K.R., Aleven, V.: Exploring the assistance dilemma in experiments with cognitive tutors. Educ. Psychol. Rev. **19**(3), 239–264 (2007). https://doi.org/10.1007/s10648-007-9049-0

Lindgren, R., Johnson-Glenberg, M.: Emboldened by embodiment: six precepts for research on embodied learning and mixed reality. Educ. Res. **42**(8), 445–452 (2013). https://doi.org/10.3102/0013189X13511661

Makransky, G., Petersen, G.B.: Investigating the process of learning with desktop virtual reality: a structural equation modeling approach. Comput. Educ. **134**, 15–30 (2019). https://doi.org/10.1016/j.compedu.2019.02.002

Mayer, R.E.: Computer games in education. Annu. Rev. Psychol. **70**, 531–549 (2019). https://doi.org/10.1146/annurev-psych-010418-102744

McLaren, B.M., Nguyen, H.A.: Digital learning games in artificial intelligence in education (AIED): a review. In: du Boulay, B., Mitrovic, A., Yacef, K. (eds.) Handbook of Artificial Intelligence in Education. Chapter 20 (2023). https://doi.org/10.4337/9781800375413.00032

McLaren, B.M., van Gog, T., Ganoe, C., Karabinos, M., Yaron, D.: The efficiency of worked examples compared to erroneous examples, tutored problem solving, and problem solving in classroom experiments. Comput. Hum. Behav. **55**, 87–99 (2016)

Nathan, M.J.: Foundations of Embodied Learning: a Paradigm for Education. Routledge (2021)

Pavlik, P.I., Jr., Brawner, K., Olney, A., Mitrovic, A.: A review of student models used in intelligent tutoring systems. In: Sottilare, R., Graesser, A., Hu, X., Holden, H. (eds.) Design Recommendations for Intelligent Tutoring Systems, Learner Modeling, vol. 1, pp. 39–68. U.S. Army Research Laboratory (2013)

Pekrun, R.: The control-value theory of achievement emotions: assumptions, corollaries, and implications for educational research and practice. Educ. Psychol. Rev. **18**(4), 315–341 (2006). https://doi.org/10.1007/s10648-006-9029-9

Renkl, A.: Toward an instructionally oriented theory of example-based learning. Cogn. Sci. **38**(1), 1–37 (2014). https://doi.org/10.1111/cogs.12086

Roschelle, J., Teasley, S.: The construction of shared knowledge in collaborative problem solving. In: O'Malley, C.E. (ed.) Computer Supported Collaborative Learning, pp. 69–97. Springer, Heidelberg (1995). https://doi.org/10.1007/978-3-642-85098-1_5

Schwartz, D.L., Bransford, J.D.: A time for telling. Cogn. Instr. **16**(4), 475–522 (1998)

Shareghi Najar, A., Mitrovic, A., McLaren, B.M.: Learning with intelligent tutors and worked examples: selecting learning activities adaptively leads to better learning outcomes than a fixed curriculum. User Model. User-Adap. Inter. **26**, 459–491 (2016). https://doi.org/10.1007/s11257-016-9181-y

Shute, V.J., Ventura, M.: Stealth Assessment: Measuring and Supporting Learning in Video Games. MIT Press. ISBN: 9780262518819 (2013)

Stahl, G., Koschmann, T., Suthers, D.: Computer-supported collaborative learning: an historical perspective. In: Sawyer, R.K. (ed.) Cambridge Handbook of the Learning Sciences, pp. 409–426. Cambridge University Press, Cambridge, MA (2006)

Van Gog, T., Kester, L., Paas, F.: Effects of worked examples, example-problem, and problem-example pairs on novices' learning. Contemp. Educ. Psychol. **36**, 212–218 (2011). https://doi.org/10.1016/j.cedpsych.2010.10.004

VanLehn, K.: The behavior of tutoring systems. Int. J. Artif. Intell. Educ. **16**(3), 227–265 (2006)

VanLehn, K.: The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educ. Psychol. **46**(4), 197–221 (2011)

VanLehn, K.: Regulative loops, step loops and task loops. Int. J. AI Educ. **26**, 107–112. (2016). https://doi.org/10.1007/s40593-015-0056-x

Yang, B.: Virtual reality and augmented reality for immersive learning: a framework of education environment design. Int. J. Emerg. Technol. Learn. (iJET). **18**, 23–36 (2023). https://doi.org/10.3991/ijet.v18i20.44209

Yaron, D., Karabinos, M., Lange, D., Greeno, J., Leinhardt, G.: The ChemCollective – virtual labs for introductory chemistry courses. Science. **328**, 584–585 (2010). https://doi.org/10.1126/science.1182435

Yusuf, H., Money, A., Daylamani-Zad, D.: Pedagogical AI conversational agents in higher education: a conceptual framework and survey of the state of the art. Educ. Technol. Res. Dev. **73**, 815–874 (2025). https://doi.org/10.1007/s11423-025-10447-4