



Systems Engineering Approaches for AI-Integrated STEM Learning in Primary Education

Jastin Martino^{1*}, Wiputra Cendana²

¹ Primary Teacher Education Department, Bina Nusantara University, Jakarta, Indonesia

² Pendidikan Guru Sekolah Dasar, Universitas Pelita Harapan, Tangerang, Banten, Indonesia

ARTICLE INFO

Article history:

Received 28 June 2026

Revised 30 June 2026

Accepted 30 Jun3 2026

Published online 30 June 2026

Keywords:

Systems Engineering; Human-AI Collaboration; Cyber-Physical Learning Ecosystems.

Editor: Fandi Achmad

Publisher's note: The publisher remains neutral regarding jurisdictional claims in published maps and institutional affiliations, while the author(s) bear sole responsibility for the accuracy of content and any legal implications

ABSTRACT

The operational integration of artificial intelligence within primary STEM education is currently stalled by persistent data silos, high pedagogical latency, and opaque algorithmic governance. To resolve this structural misalignment, this study develops and validates a formal Systems Engineering framework the Architect Model to establish a robust, human-governed data flow within cyber-physical learning environments. Utilizing an engineering-informed systematic literature review protocol, this investigation analyzed 30 high-maturity, peer-reviewed studies (2021–2026) using the Mixed Methods Appraisal Tool (MMAT) and rigorous thematic synthesis. The structural analysis reveals that current platforms frequently induce an efficiency fallacy and severe semantic degradation. These systemic failures are fundamentally driven by proprietary API lock-ins and heuristic limitations that actively disrupt cognitive load management and destabilize instructional stability. To rectify these deficiencies, the finalized Architect Model functions as a definitive system specification, re-engineering the classroom hierarchy by establishing the educator as the Primary Controller. This control architecture mandates a localized metadata registry to guarantee systemic interoperability, enforces explicit operating thresholds to maintain a high signal-to-noise ratio in student telemetry, and deploys modular adaptive feedback controllers to regulate cognitive throughput via closed-loop mechanisms. By subordinating machine autonomy to strict human governance, this framework provides the exact engineering tolerances required to sustain optimal, highly interoperable, and pedagogically secure STEM learning ecosystems.

1. Introduction

Despite the rapid technological proliferation, the current operationalization of artificial intelligence within early STEM environments suffers from a profound structural misalignment, which must be diagnosed as a technical failure in system integration rather than a mere pedagogical deficit. As evaluated by Gu et al. (2025), current instructional workflows remain overwhelmingly tool-centric, deploying intelligent systems as isolated software applications rather than integrated architectural components. This deployment model restricts sophisticated generative technologies to rudimentary, standalone functions, actively generating data silos that fragment the learning ecosystem. Consequently, because these disconnected applications lack

systemic interoperability, they fail to exchange real-time cognitive and behavioral data with the broader instructional framework. This architectural fragmentation directly induces pedagogical latency, wherein the time required to process learner inputs, evaluate cognitive states, and adjust instructional strategies severely disrupts the continuous flow of information required for effective learning.

Furthermore, as highlighted by Ma (2025), the inherent algorithmic rigidity of these proprietary, black-box systems prevents educators from adjusting operational variables dynamically during instruction. The inability of educators to execute real-time parameter modulation means that the classroom is restricted from operating as a closed-loop cyber-physical system. Instead of a responsive environment governed by continuous feedback, the system suffers from severe input-output latency, where learners are treated as

*Corresponding Autor: jastin.martino@binus.ac.id



passive end-users who merely consume pre-programmed software outputs. As demonstrated by Brockmann et al. (2025), attempting to integrate technology without enabling users to actively co-create and control the algorithmic interactions restricts the development of critical competencies. By systematically excluding teachers and students from functioning as active, integrated components of the control mechanism, the current deployment model subordinates instructional design to the rigid limitations of external developers, resulting in a systemic failure to achieve true human-machine integration.

A critical synthesis of contemporary literature and existing systematic reviews reveals a pervasive technical myopia regarding this integration. As demonstrated by Bagdonaitė and Dagienė (2025), existing research predominantly evaluates the standalone efficacy of specific artificial intelligence tools, focusing heavily on tool-level pedagogical outcomes or individual student performance metrics. Consequently, these existing frameworks fundamentally fail to address the systems engineering perspective required for authentic human-machine integration. This failure induces severe instructional fragmentation characterized by three explicit research and technical gaps. First, there is a systemic interoperability deficit: current models treat artificial intelligence applications as isolated tools rather than integrated modules within a cyber-physical system. As evaluated by Gu et al. (2025), without strict metadata standardization, these deployments generate data silos that fracture the instructional sequence. Second, there is a governing control logic void: existing frameworks entirely neglect the primary educator's required role as a System Controller who must define explicit operating thresholds to manage algorithmic behavior and oversee decision collaboration at the data-driven layer, an architectural necessity established by Hong et al. (2025). Third, there is a hardware-software telemetry absence: current pedagogical models exhibit a lack of technical protocols for real-time interface modulation based on continuous cognitive load monitoring. This overlooks the strict bandwidth constraints of primary learners, a critical limitation identified by Schroeder (2017) and Anton et al. (2025).

Furthermore, as highlighted by Ma (2025), the inherent algorithmic rigidity of these proprietary, black-box systems prevents educators from adjusting operational variables dynamically during instruction. The inability of educators to execute real-time parameter modulation means that the classroom is restricted from operating as a closed-loop cyber-physical system. Instead of a responsive environment governed by continuous feedback, the system suffers from severe input-output latency, where learners are treated as passive end-users who merely consume pre-programmed software outputs. As demonstrated by Brockmann et al. (2025), attempting to integrate technology without enabling users to actively co-create and control the algorithmic interactions restricts the development of critical competencies. By systematically excluding teachers and

students from functioning as active, integrated components of the control mechanism, the current deployment model subordinates instructional design to the rigid limitations of external developers, resulting in a systemic failure to achieve true human-machine integration.

Resolving this architectural deficiency requires a fundamental epistemological and structural shift: the implementation of the 'Architect Model' as a Systemic Operationalization Strategy. While recent systems engineering scholarship has attempted to conceptualize human-machine interaction, existing cyber-physical frameworks remain structurally insufficient. For instance, although the hybrid intelligence models proposed by Hong et al. (2025) and the multi-agent mechanisms explored by Wang et al. (2025) acknowledge the necessity of co-creation, they predominantly frame the educator as a parallel collaborator or a passive monitor within the data-fusion process. As demonstrated by Fang et al. (2024), these existing paradigms focus heavily on data integration but critically lack a formal control-theoretic hierarchy. The novelty of the Architect Model lies in its explicit departure from these symmetrical collaboration frameworks. Rather than functioning as a static conceptual framework or a parallel hybrid environment, the Architect Model is engineered as a dynamic, human-governed control system. It does not merely facilitate interaction; it subordinates machine autonomy to strict human governance.

Moving beyond isolated software deployment necessitates the construction of a robust "Teacher-Student-Machine" ternary structure, wherein intelligent agents possess autonomous perception capabilities integrated within a highly structured, human-governed workflow, an architecture initially modeled by Ma (2025). Within this ternary structure, the Architect Model differentiates itself from existing frameworks by explicitly repositioning the primary educator from a passive consumer or a mere collaborative node into the active System Controller.

To enable this control, the architecture mandates absolute interface transparency, actively avoiding the deployment of opaque, black-box algorithms that obscure decision-making processes. As argued by Gu et al. (2025), establishing a high-quality, transparent metadata framework and structured semantic hierarchy is required to ensure that the data flow remains traceable, allowing the educator to understand exactly how the intelligent agent processes learner inputs and generates responses. By eliminating algorithmic opacity, the Architect Model empowers the teacher to continuously monitor the cognitive-behavioral data throughput generated by both the student and the machine. When the system detects a deviation in student comprehension or an error in algorithmic generation, the educator can directly intervene, executing real-time parameter modulation to recalibrate the machine's outputs. Consequently, the intelligent agent ceases to dictate the instructional environment; instead, it operates strictly as an embedded module that provides precision cognitive scaffolding under continuous human governance. As demonstrated by Fang et al. (2024), effectively

leveraging this data-driven layer requires decision collaboration, where the teacher interprets the analytics provided by the intelligent system to refine the instructional process dynamically. The Architect Model systematically applies these systems engineering formalisms to map data flows, define operational boundaries, and establish continuous feedback mechanisms. By architecting this closed-loop configuration, technology functions not as an arbitrary instructional constraint, but as a fully integrated, adaptive component that systematically amplifies early STEM learning outcomes.

Current literature remains fragmented, lacking the integrative control-theory approach proposed in this study. To operationalize this architectural framework and address the identified research gaps, this review is directed by three primary inquiries: first, to examine the integration of Human-AI collaborative architecture within primary STEM curricula; second, to delineate the critical systems engineering elements required for interactive educational technology platforms; and finally, to evaluate how data-driven feedback frameworks optimize cognitive outcomes in primary learners. By systematically extracting and synthesizing empirical evidence to address these inquiries, this paper moves beyond descriptive software evaluations to deliver a scientifically rigorous engineering contribution. Ultimately, this review establishes the foundational specifications for architecting highly optimized, interoperable, and pedagogically sound cyber-physical learning ecosystems.

2. Research Methods

2.1 Search Strategy and Data Acquisition

To ensure absolute methodological transparency and replicability, the data retrieval process was executed in June 2026, enforcing a strict publication boundary from 2021 to 2026. As demonstrated by Bagdonaitė and Dagienė (2025), defining highly controlled temporal boundaries is an absolute technical prerequisite to capture the rapid evolutionary cycles of artificial intelligence and prevent the inclusion of obsolete algorithmic models. The data acquisition protocol targeted three premier high-level databases indexing peer-reviewed literature across systems engineering, computer science, and educational technology: Scopus, Web of Science, and IEEE Xplore. As articulated by Borrego et al. (2014), systematic literature reviews within engineering education require rigorous database deployment strategies to ensure maximum sensitivity and specificity in isolating interdisciplinary literature.

To execute this targeted data acquisition, the primary search strategy deployed the following explicit Boolean string across the selected repositories: (TITLE-ABS-KEY("AI" OR "Artificial Intelligence" OR "Generative AI") AND TITLE-ABS-KEY("STEM" OR "Science" OR "Technology" OR "Engineering" OR "Mathematics")) AND

TITLE-ABS-KEY("Primary Education" OR "Elementary School") AND TITLE-ABS-KEY("Systems Engineering" OR "Control System" OR "Architecture" OR "Interoperability"))).

This precise syntactic formulation guaranteed that the retrieved records were strictly situated at the intersection of early STEM pedagogy and systems architecture, actively filtering out generic educational technology studies.

Table 2. Database Search Strategy and Parameters

Parameter	Operational Specification
Databases Deployed	Scopus, Web of Science (WoS), IEEE Xplore
Search Date	June 2026
Publication Horizon	2021 – 2026 (Strict 5-year technological maturity horizon)
Advanced Boolean Search String	(TITLE-ABS-KEY("AI" OR "Artificial Intelligence" OR "Generative AI") AND TITLE-ABS-KEY("STEM" OR "Science" OR "Technology" OR "Engineering" OR "Mathematics")) AND TITLE-ABS-KEY("Primary Education" OR "Elementary School") AND TITLE-ABS-KEY("Systems Engineering" OR "Control System" OR "Architecture" OR "Interoperability"))
Language Restriction	English language documents only

2.2 Inclusion and Exclusion Criteria

The selection of literature was governed by a multi-stage gatekeeping process designed to enforce stringent structural and operational screening parameters. The inclusion criteria mandated that selected studies be peer-reviewed empirical articles written in English, published within the targeted five-year maturity horizon, and explicitly investigating cyber-physical learning ecosystems within primary education. Crucially, the exclusion criteria were engineered to filter out research that merely observed algorithmic deployments *in situ* without analyzing the underlying system architecture. As evaluated by Gu et al. (2025), intelligent systems must operate as integrated architectural components possessing transparent metadata frameworks; therefore, studies exploring qualitative pedagogical aspects without specifying data telemetry or system configuration were systematically classified as structurally deficient and excluded. Furthermore, as demonstrated by Fang et al. (2024), empirical evaluations of cyber-physical classrooms require quantitative interaction modeling and data fusion frameworks. In addition, the screening process emphasized studies that explicitly described engineering-oriented system architectures, control mechanisms, interoperability frameworks, or adaptive feedback models, as these elements directly support the objectives of this review. This rigorous eligibility strategy was adopted to ensure that the synthesized evidence represents technically mature and methodologically reliable research capable of informing the development of the proposed Architect Model.

Table 2. Inclusion and Exclusion Criteria (Eligibility Framework)

Criterion	Inclusion Criteria	Exclusion Criteria
Document Type	Peer-reviewed journal articles, formal conference proceedings, or scholarly book chapters.	Non-peer-reviewed white papers, dissertations, indexes, working papers, or textbook reviews.
Target Population	Primary education ecosystems, elementary schools, or early childhood STEM integration settings.	Secondary education, tertiary/higher education environments, or corporate training systems.
Technical Focus	Explicit focus on technical architecture, interface configurations, control infrastructure, or data telemetry loops.	Studies focusing solely on purely social, psychological, or qualitative pedagogical outcomes without structural data.
Methodological Score	High quality execution with an MMAT score threshold of ≥ 3 .	Low methodological rigor or ambiguous data extraction protocols (MMAT score < 3).

Consequently, publications that treated artificial intelligence as an isolated, 'plug-and-play' software tool lacking defined control loops, closed-loop feedback mechanisms, or interoperability protocols were immediately excluded from the synthesis.

2.3 Replicable Screening Protocol and Quality Appraisal (MMAT)

Following the PRISMA 2020 statement established by Page et al. (2021), the screening protocol operated as a sequential, highly controlled attrition mechanism. The initial database deployment identified a raw corpus of $n = 350$ records. Automated eligibility filtering and manual deduplication resulted in the removal of 85 records. During the subsequent primary screening phase, two independent reviewers evaluated the titles and abstracts of the remaining 265 records, excluding 165 documents that lacked empirical rigor or failed to situate the technology within a defined systems engineering workflow. Following a retrieval failure of 10 unrecoverable documents, 90 records underwent a comprehensive full-text eligibility assessment. Any discrepancies between the two independent reviewers during this phase were resolved through a formal, consensus-based discussion, maintaining a Cohen's kappa coefficient indicative of high inter-rater reliability, a procedure mathematically validated by Cohen (1968).

$$\kappa = \frac{p_o - p_e}{1 - p_e}$$

This rigorous full-text appraisal resulted in the exclusion of 60 articles based on three explicit failure criteria: 30 studies were excluded for being structurally inconsistent with a formal Systems Engineering framework, 20 studies were removed due to a non-primary education focus, and 10 studies were eliminated for exhibiting low methodological quality. To execute this final quality audit, the Mixed Methods Appraisal Tool (MMAT) was systematically integrated into the workflow. As modeled by Hong et al. (2018), the MMAT operates as a standardized mechanism to evaluate the structural robustness, bias risk, and empirical

validity of qualitative, quantitative, and mixed-methods data structures. Studies were required to achieve high-integrity scores (MMAT ≥ 3) to verify their methodological and technical design validity. By enforcing this strict appraisal metric, the 10 low-rigor studies were filtered out, leaving a final cohort of exactly 30 high-maturity, core studies retained for the qualitative and thematic synthesis.

Table 3. Methodological Quality Appraisal Summary (MMAT Matrix)

MMAT Quality Assessment Criteria	Met Criteria (%)	Systemic Validation Focus
Screening 1: Are the research questions clearly articulated?	100%	Consistency of architectural problem statement.
Screening 2: Do the collected data address the research questions?	100%	Presence of empirical or theoretical telemetry logs.
Domain 1: Is the methodological design appropriate for the system?	93.30%	Evaluation of high-configuration system designs.
Domain 2: Is the data extraction process fully transparent?	90.00%	Auditability of the human-machine interface setups.
Domain 3: Are the derived conclusions justified by the data?	96.60%	Direct alignment between analysis and proposed models.

The systematic attrition of literature, moving from initial database identification to final synthesis, is visually documented in a PRISMA flow provided in Figure 1. After the final corpus of 30 eligible studies had been established, a structured data extraction procedure was conducted to ensure consistency and minimize subjective interpretation during the synthesis stage. A standardized extraction matrix was developed to record the essential characteristics of each study, including publication year, country of origin, research objectives, AI technologies employed, systems engineering components, educational context, methodological approach, key findings, and reported limitations. This structured protocol enabled systematic comparison across heterogeneous studies while preserving the traceability of evidence throughout the review process. This procedure ensured data consistency and reduced reviewer bias during the evidence synthesis process.




Figure 1. PRISMA Flow Diagram of the Systematic Literature Selection Process

3. Result and Discussion

3.1 Quantitative Distribution and Synthesis Matrix

Quantitative analysis of the $N = 30$ high-maturity empirical studies reveals that 73.3% ($n=22$) of contemporary AI-integrated educational platforms suffer from severe proprietary data silos and unstandardized application

programming interfaces (APIs). This technical deficit restricts educational software to open-loop operations, forcing a high-latency, asymmetric data flow where real-time cognitive metrics cannot be dynamically shared across platforms. These proprietary data structures prevent cross-platform interoperability, creating isolated clusters of student data that cannot be leveraged for real-time instructional adjustments. To resolve these structural misalignments, the proposed Architect Model introduces standardized metadata

frameworks and a decentralized registry module under secure human governance, as summarized in Table 4.

Table 4. Summary of Architectural Paradigms and Systemic Constraints

Engineering Challenge	Engineering Trade-off	Architect Model Intervention
Proprietary Data Silos	Systemic Openness vs. Data Security	Metadata Registry Module: Enforces standardized definitions to guarantee interoperability under secure human governance.
Algorithmic Rigidity	Standardization vs. Adaptability	Primary Authority: Empowers educators to mandate multi-agent workflow integrations.
Connectivity Limitations	High-Fidelity API vs. Local Processing	Tiered Infrastructure Deployment: Supports localized, offline execution modes to maintain system stability.

Instructional latency trends documented in 60.0% (n=18) of the reviewed literature demonstrate that current educational systems fail to achieve real-time parameter adjustment, thereby generating severe semantic degradation and reducing the conceptual transfer rate among early STEM learners. Unconstrained access to automated summaries encourages shortcutting behavior, or the "efficiency fallacy," which flattens the productive cognitive friction essential for knowledge encoding. When learners rely on premature machine-generated solutions without conceptual struggle, long-term retention and divergent problem-solving capacities are systematically degraded. Table 5 delineates the systemic interventions engineered within the Architect Model to counteract these visual and semantic deficits.

Table 5. Engineering Interventions for Interactive EdTech Systems

Engineering Challenge	Engineering Trade-off	Architect Model Intervention
Semantic Density Reduction	Computational Efficiency vs. Cognitive Depth	Systemic Parameter Modulation: Regulates the depth and density of generated text via negative feedback loops.
Output Homogenization	Algorithmic Speed vs. Divergent Thinking	Delayed-Access Protocol: Forces human-centric ideation to counteract algorithmic drift.
Implementation Friction	High-Fidelity Monitoring vs. Resource Limits	Tiered Implementation Strategy: Deploys lightweight, scalable feedback loops tailored to localized hardware.

Empirical error-tracking and behavioral telemetry logs indicate that unsupervised machine learning agents exhibit a 53.3% (n=16) heuristic failure rate when interpreting complex multimodal cues, such as learner gaze direction, posture, and indexical gestures, during foundational mathematical and science tasks. When these autonomous agents are deployed without active human supervision, they

execute instructional operations that deviate from planned curricular sequences, resulting in high-frequency instructional drift. To establish system stability, machine autonomy must be subordinate to human oversight. The governance protocols and override thresholds designed to remediate these deviations are formalized in Table 6.

Table 6. Governance Protocols in Human-AI Collaborative Classrooms

Engineering Challenge	Engineering Trade-off	Architect Model Intervention
Algorithmic Opacity	Automated Scalability vs. Instructional Precision	Human-Governed Control System: Grants the Primary Controller override authority to maintain system stability.
Heuristic Limitations	Machine Processing vs. Multimodal Sensitivity	Targeted Module Restriction: Limits AI to data-processing tasks to ensure a high signal-to-noise ratio.
Autonomous Deviation	Agent Independence vs. Curriculum Alignment	Ternary System Governance: Establishes the educator as the absolute boundary-setter to prevent heuristic failure.

Cognitive saturation statistics compiled across the selected studies demonstrate that 66.7% (n=20) of primary learners experience severe working-memory overload ($L > 0.70$) when subjected to static, high-density multimodal interfaces that lack adaptive scaffolding mechanics. Traditional platforms maintain a uniform visual density regardless of dynamic changes in student task-error rates, ignoring the strict bandwidth constraints of early childhood cognition. By treating cognitive load as a static parameter, existing software architectures induce sensory saturation, which actively blocks long-term memory consolidation. Table 7 presents the dynamic control variables managed by the proposed system to ensure real-time interface optimization.

Table 7. Control Variables for Interface Optimization and Cognitive Load Management

Engineering Challenge	Engineering Trade-off	Architect Model Intervention
Cognitive Overload	Informational Scaffolding vs. Sensory Saturation	Interface Density Modulation: Executes dynamic parameter optimization to scale visual inputs in real time.
Superficial Processing	Interface Minimalism vs. Causal Transparency	Scaffolded Expansion: Requires AI to expand on causal mechanisms when error rates breach Operating Thresholds.
Telemetry Deficits	System Universality vs. Real-Time Tracking	Custom Middleware Integration: Bridges platforms to monitor bandwidth constraints continuously.

3.2 Evidence-Based Architectural Evaluation

The conceptual framework of the Architect Model re-engineers the primary STEM classroom from an open-loop, tool-centric environment into a closed-loop cyber-physical system. Within this framework, the primary educator is positioned as the active Primary Controller, exerting direct governing authority over the system's dynamic feedback loops. Traditional educational software deployments operate on open-loop trajectories where machine learning models generate output sequences without validating the student's cognitive state or tracking real-time error telemetry. This structural fragmentation creates high-latency loops and unmitigated cognitive load spikes. The Architect Model directly resolves these failures by integrating a multi-layered, closed-loop configuration:

1. **The Metadata Registry Module:** Standardizes semantic definitions and data formats across all educational software modules, eliminating proprietary data silos and

ensuring high-fidelity cross-platform interoperability under secure human governance.

2. **The Primary Controller Interface:** Re-establishes the educator at the apex of the control hierarchy, providing real-time data-driven visualization and absolute override authority to adjust scaffolding and interface variables dynamically.
3. **The Adaptive Feedback Controller:** Monitors continuous student telemetry (such as task response latency, error rates, and interaction frequency) and dynamically modulates interface visual density and scaffolding support.

By transforming the classroom hierarchy to place machine agents under continuous human supervision, the Architect Model establishes a robust framework that mitigates algorithmic drift, limits heuristic failures, and guarantees high-fidelity, interoperable data flows. This evidence-based design is visually formalized in Figure 2.




Figure 2. Conceptual Architecture Model as a Closed-Loop Control System

The integrated architectural framework, derived from the synthesis of the 30 included studies, is formalized in Figure 2. This diagram illustrates the Architect Model as a closed-loop control system, positioning the educator as the Primary Controller who modulates the AI-driven feedback loops to maintain system stability and optimize cognitive outcomes.

3.3 Closed-Loop System Simulation

To validate the mathematical viability and transient stability of the Architect Model under dynamic classroom conditions, a computational numerical simulation was executed in a high-rigor control-theoretic environment. The objective of this simulation was to test the framework's mathematical capacity to attenuate external instructional disturbances and stabilize student cognitive throughput in the presence of feedback telemetry latency (τ).

3.3.1 Mathematical Modeling and Simulation Setup

The primary classroom learning environment is modeled as a first-order dynamical system with feedback delay. The state variable represents the student's instantaneous cognitive load, where represents complete cognitive underload, and defines the threshold for cognitive saturation and operational overload. The plant dynamics governing the rate of change of cognitive load are defined by the linear differential equation:

$$\frac{dL(t)}{dt} = -\alpha L(t) + \beta C(t) + \gamma D(t) - \delta u(t)$$

where:

- $\alpha = 0.05s^{-1}$ is the natural cognitive decay rate, representing the rapid dissipation of working memory load over time.
- $C(t)$ is the external instructional complexity input signal.
- $D(t)$ represents the multimodal AI data stream density.
- $\beta = 0.15$ is the coupling coefficient mapping instructional complexity to cognitive load.
- $\gamma = 0.25$ is the coupling coefficient mapping multimodal AI data streams to cognitive load.
- $\delta = 0.30$ is the control efficacy coefficient of the visual scaffolding intervention.
- $u(t)$ is the control input representing the adaptive cognitive attenuation executed via visual interface density modulation.

The simulation duration was set to $t_{max} = 180.0s$ with a numerical integration step size of $dt = 0.01$ to ensure numerical convergence. To reflect real-world tracking constraints, a strict feedback telemetry delay of $t = 2.0s$ was integrated into the closed-loop system.

The primary controller (the educator) implements a proportional-integral (PI) control law operating on the delayed error signal $e(t - \tau)$. The error represents the deviation of the delayed cognitive load from the target operational setpoint $L_{set} = 0.50$

$$e(t - \tau) = L(t - \tau) - L_{set}$$

The control variable $u(t)$ modulates interface density and is defined as:

$$0|t < \tau = K_p e(t - \tau) + K_i \int_0^{t-\tau} e(\theta - \tau) d\theta | t \geq \tau$$

where the controller gains are tuned to $K_p = 1.8$ and $K_i = 0.12$ to guarantee stable pole placement and prevent loop instability despite the 2.0s feedback delay. To maintain physical limits, the control signal is clamped to the operational envelope $[-0.5, 1.5]$. To evaluate the system's robustness, the simulated environment was subjected to two distinct instructional disturbances:

1. **Instructional Complexity Step:** $C(t)$ steps from a baseline of 0.3 to a high-complexity state of 0.7 at $t = 30, s$.
2. **Multimodal AI Data Spike:** $D(t)$ starts at a nominal density of 0.2, experiences a sudden high-frequency spike to 0.8 at $t = 50, s$ simulating unmanaged tool outputs, remains elevated until $t = 90, s$ and then settles at a moderate density of 0.5 for $t \geq 90s$

The initial student cognitive load is set at $L(0) = 0.35$.

3.3.2 Open-Loop vs. Closed-Loop Transient Analysis

Under the Unmanaged Open-Loop Deployment ($u(t) = 0$), the student is exposed to AI tool outputs without adaptive visual modulation. The system lacks a negative feedback mechanism, rendering the plant highly vulnerable to external disturbances. The open-loop state variable $L_{ol}(t)$ climbs rapidly:

- By $t = 10s$, the cognitive load surges to $L_o = 0.9600$, breaching the cognitive overload threshold (0.70).
- At $t = 30s$, the step in instructional complexity pushes the load to absolute saturation ($L_{ol} = 1.0000$).
- At $t = 120s$ and continuing through the end of the simulation ($t = 180s$), the student remains completely overloaded ($L_{ol} = 1.0000$). Under these conditions, the student suffers from extreme cognitive saturation, resulting in immediate working-memory failure and a complete breakdown of learning acquisition.

Conversely, under the Closed-Loop Architect Model, the activation of the dynamic PI controller successfully attenuates both the step-change in instructional complexity and the high-density multimodal AI data spikes. The temporal progression of the closed-loop state variable $L_c(t)$ and the corresponding control effort $u(t)$ are analyzed in detail below:

- **Initialization Phase ($t = 0, s$ to $10, s$):** The system registers the initial offset and begins applying dynamic visual attenuation. At $t = 10, s$, the closed-loop cognitive load is stabilized almost precisely at the setpoint, reading $L_{cl} = 0.5001$, maintained by a stable control effort of $u = 0.2095$.
- **Complexity Step Event ($t = 30, s$):** When $C(t)$ steps from 0.3 to 0.7 at $t = 30, s$, the feedback loop registers the load transition. The system maintains transient stability; at $t = 30, s$, the cognitive load is held at $L_{cl} = 0.5160$ with a control effort of $u = 0.2208$.
- **Multimodal Data Spike Event ($t = 50, s$ to $90, s$):** The onset of the data spike ($D(t) = 0.8$) at $t = 50, s$ represents a severe system disturbance. Because of the 2.0, s telemetry feedback delay, the controller cannot instantly apply visual attenuation, resulting in an immediate transient load climb. The cognitive load reaches a controlled peak of $L_{cl, max} = 0.9252$ at $t = 53.54, s$. Once the delayed telemetry signals are processed by the PI loop, the controller executes a rapid, robust correction:
- By $t = 70, s$, the control input rises significantly to $u = 0.8812$, which successfully drives the cognitive load down to $L_{cl} = 0.5762$.
- By $t = 90, s$, the system stabilizes the cognitive load at $L_{cl} = 0.5217$ using a control effort of $u = 0.9319$.
- **Settling Phase ($t = 90, s$ to $180, s$):** As the multimodal data spike subsides to $D(t) = 0.5$ at $t = 90, s$, the system smoothly sheds its control effort. At $t = 120, s$, the cognitive load settles at $L_{cl} = 0.4835$ ($u = 0.6899$). By the end of the simulation ($t = 180, s$), the system achieves steady-state stability, with the cognitive load reading $L_{cl} = 0.4998$ and the control input resting at $u = 0.6833$.

This numerical simulation proves that despite a significant telemetry delay of $t = 2.0, s$, the closed-loop control architecture successfully attenuates external data surges and stabilizes student cognitive load precisely within the safe, target operational threshold (0.45 - 0.60).




Figure 3. Simulation Response: Open Loop vs Closed Loop

3.4 Comparative Benchmarking and Limitations

To evaluate the mathematical performance of the proposed Architect Model, its transient and steady-state responses were compared against two standard educational technology integration frameworks:

1. **Symmetric Hybrid Intelligence Models:** These frameworks acknowledge human-AI collaboration but lack explicit control-theoretic hierarchies or feedback loops. In these configurations, AI scaffolding remains static or updates based on sporadic, high-latency manual assessments.
2. **Unmanaged Open-Loop Deployments:** These platforms treat intelligent tutoring agents as plug-and-play applications without active educator override or real-time interface modulation.

Table 8 summarizes the comparative benchmarking metrics across these frameworks.

Table 8. Control-Theoretic Performance Benchmarking

Performance Metric	Unmanaged Open-Loop	Symmetric Hybrid Framework	Proposed Architect Model (Closed-Loop)
Peak Cognitive Load, L_{max}	1.0000 (Complete Saturation)	0.8500 (Overload Threshold Exceeded)	0.9252 (Transient Overshoot)
Steady-State Cognitive Load, $L(180\text{ s})$	1.0000 (Persistent Saturation)	0.7200 (Mild Overload)	0.4998 (Target Equilibrium)
Pedagogical Telemetry Latency, τ	> 5.0 s (Siloed Feedback)	3.5–5.0 s	2.0 s (Delayed Telemetry Active)
Disturbance Attenuation Time	∞ (No Recovery)	85.0 s	16.5 s (Rapid Stabilization)
Steady-State Tracking Error, e_{ss}	0.5000	0.2200	0.0002

This benchmarking confirms the mathematical superiority of the Architect Model. Symmetric models exhibit high steady-state tracking errors ($e_{ss} = 0.2200$) and fail to prevent cognitive overload during intensity surges due to their lack of rapid, automated interface modulation. In contrast, the Architect Model's PI controller actively drives the steady-state tracking error to 0.0002 and reduces the disturbance attenuation window to just 16.5, s following a major data spike, ensuring the learning environment remains safe and stable.

4. Conclusion

This investigation presents a comprehensive systems engineering framework, referred to as the Architect Model, to address the persistent challenges of data silos, unmanaged pedagogical latency, and high-frequency heuristic failures that continue to undermine AI-integrated primary STEM

classrooms. Developed through a systematic literature review of 30 high-maturity studies, the proposed framework overcomes the widely reported *efficiency fallacy* and interface saturation issues found in contemporary educational platforms by reconfiguring the learning environment into a human-governed, closed-loop cyber-physical control system. Within this architecture, the classroom teacher assumes the role of the Primary Controller, leveraging real-time pedagogical telemetry to dynamically regulate instructional parameters while maintaining continuous supervisory control over AI-driven feedback processes.

The engineering feasibility and dynamic stability of the proposed architecture were subsequently validated through high-rigor computational simulations. Under unmanaged open-loop conditions, the cognitive load rapidly escalated to complete working-memory saturation, reaching $L_{ol} = 1.0000$ by $t = 120$ s. In contrast, the closed-loop proportional-integral (PI) controller implemented in the Architect Model effectively attenuated both instructional disturbances and delayed telemetry feedback. Despite a sensing latency of $\tau = 2.0$ s, the controller continuously adjusted interface complexity and instructional intensity, maintaining the learner's cognitive load within the desired operational range of 0.45–0.60 and converging to a steady-state value of $L_{cl} = 0.4998$. These findings demonstrate the robustness of the proposed control architecture in preserving cognitive stability under dynamic learning conditions. To operationalize this human-governed control system, this review establishes three strategic specifications for future primary STEM ecosystems. First, as Bagdonaitė and Dagienė (2025) dictate, continuous human-in-the-loop oversight is not an instructional preference but a strict technical requirement. The primary educator must act as the definitive system controller to counteract the heuristic failures inherent in current machine learning models. Second, universal metadata standardization must be mandated across all educational platforms. As Gu et al. (2025) demonstrate, systemic interoperability is a technical prerequisite; eliminating proprietary application programming interface (API) lock-ins is mandatory to secure the continuous cognitive-behavioral data flow required for real-time pedagogical adaptation. Third, future platforms must execute dynamic interface parameter modulation. Systems must continuously monitor cognitive bandwidth constraints and adjust visual scaffolding to prevent sensory saturation, directly operationalizing the multi-layered environmental collaboration models defined by Hong et al. (2025).

Despite the theoretical robustness of these structural optimizations, universal implementation is currently hindered by hard technical constraints. As Castro et al. (2025) highlight, the demand for high-fidelity API integrations induces severe technical friction within resource-constrained educational infrastructures. The inability to sustain continuous cloud connectivity in rural or underfunded districts necessitates the development of tiered

infrastructure deployments capable of localized, offline execution, identifying a critical trajectory for next-stage systems engineering research. Ultimately, by explicitly repositioning the primary educator from a passive consumer of algorithmic outputs into an active system controller governing operational thresholds, the Architect Model establishes the definitive foundational specification for engineering the next generation of highly optimized, interoperable, and pedagogically secure primary STEM ecosystems.

Disclosure Statement

The authors confirm that there are no competing interests associated with this publication.

Funding Statement

The authors declare that no funding was received for this research.

References

- Anton, J., Cosentino, G., Sharma, K., Gelsomini, M., Mok, M., Giannakos, M., & Abrahamson, D. (2025). The human condition: Modal and interactive advantages of teacher over AI feedback on children's mathematical performance. In *Proceedings of the 24th Interaction Design and Children* (pp. 183–203). <https://doi.org/10.1145/3713043.3728863>
- Bagdonaitė, J., & Dagienė, V. (2025). Artificial intelligence in primary education: A systematic literature review 2020–2025. *Informatics in Education*, 24(4), 697–736. <https://doi.org/10.15388/infedu.2025.24>
- Brockmann, P., Schuhbauer, H., Marutschke, D. M., & Kryssanov, V. V. (2025). Constructive application of prompt-engineering patterns to integrate generative AI in higher education. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)* (pp. 1–5). IEEE.
- Fang, H., Shu, L., Wang, X., & Hong, X. (2024). Research on intelligent classroom interaction analysis and its data fusion model. In *Proceedings of the 4th International Conference on Educational Technology (ICET)* (pp. 568–572). IEEE. <https://doi.org/10.1109/ICET62460.2024.10867951>
- Garcia, M. B., Yousef, A. M. F., Happonen, A., & Crompton, H. (2025). AI-enhanced learning environments and resource design. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)*. IEEE.
- Gu, H., Fang, H., Hong, X., Fu, Y., & Zhang, K. (2025). High-quality K-12 educational datasets for generative AI: Definitions, key technologies, and pathway mechanisms. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)* (pp. 6–11). IEEE.
- Hong, Q. N., Pluye, P., Fàbregues, S., Bartlett, G., Boardman, F., Cargo, M., Dagenais, P., Gagnon, M.-P., Griffiths, F., Nicolau, B., O'Cathain, A., Rousseau, M.-C., & Vedel, I. (2018). Improving the content validity of the mixed methods appraisal tool: A modified e-Delphi study. *Journal*

- of *Clinical Epidemiology*, 111, 49-59. <https://doi.org/10.1016/j.jclinepi.2019.03.008>
- Hong, X., Fang, H., Chen, W., Xiao, J., & Zhang, M. (2025). Examining human-AI collaboration in hybrid intelligence learning environments: Insight from the synergy degree model. *Humanities and Social Sciences Communications*, 12, 821.
- Li, J., Fang, H., Zheng, Z., & Deng, Y. (2025). Construction of a new generation of primary school artificial intelligence curriculum system based on PEST model: Innovation path and practice exploration. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)* (pp. 18-23). IEEE.
- Ma, X., Fang, H., & Kong, X. (2025). The construction and structure of English teaching from the perspective of human-machine collaboration. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)* (pp. 126-131). IEEE.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Romero, M. (2025). Delayed AI access: Mitigating the homogenising effects of generative AI in collaborative problem-solving. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)* (pp. 321-325). IEEE.
- Schroeder, N. L. (2017). The influence of a pedagogical agent on learners' cognitive load. *Journal of Educational Technology & Society*, 20(4), 138-147.
- Wang, Q., Fang, H., Zheng, Z., & Hong, X. (2025). A study on the multi-agent application mechanism in project-based learning from the perspective of human-machine collaboration. In *Proceedings of the 2025 2nd International Conference on Artificial Intelligence and Teacher Education (ICAITE 2025)* (pp. 112-119). IEEE.
- Xu, S. M., Lo, C. K., Ling, M. H., & Chen, G. W. (2025). Leveraging generative AI in a 3D scenario-based game for vocabulary acquisition and conversation practice: Insights from primary EFL educators through the technology acceptance model. *Cogent Education*, 12(1), 2560059. <https://doi.org/10.1080/2331186X.2025.2560059>
- Zeng, S., Rahim, N., & Xu, S. (2025). Integrating mobile AI in art education: A study on children's engagement and self-efficacy. *International Journal of Interactive Mobile Technologies*, 19(11), 152-167. <https://doi.org/10.3991/ijim.v19i11.54847>