AI-Driven Personalized Learning in Educational Systems: A Framework for Adaptive Learning and Decision Support

Vishaliney Pirathap[1], Sangeetha Arunprakash[2], M. S. S. Mayadunna[3]

Department of Information Technology, Faculty of Computing, SLIIT Northern uni, Jaffna, Sri Lanka

***Corresponding Author**

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150700094

Received: 29 July 2026; Accepted: 03 Auguust 2026; Published: 15 Auguust 2026

ABSTRACT

Artificial intelligence (AI) has expanded the capacity of educational systems to support learner modelling, performance prediction, adaptive resource recommendation, automated feedback, and conversational tutoring. However, existing AI-based personalized learning systems commonly employ these capabilities as separate technical components, with limited coordination, pedagogical grounding, explainability, educator control, and governance. This study proposes the Pedagogically Constrained Explainable Orchestration (PCEO) framework to address these limitations. A structured literature review guided by PRISMA 2020 reporting principles and critical thematic synthesis was used to identify recurring technical, educational, ethical, and human-oversight requirements in AI-driven personalized learning. These findings were translated into a requirement-driven conceptual architecture. The proposed framework coordinates specialized agents for learner modelling, engagement and risk analysis, recommendation, policy optimisation, generative support, and governance through a shared learner-state representation and an explicit orchestration mechanism. Candidate educational actions are evaluated using pedagogical suitability, expected learning benefit, engagement, explainability, cognitive load, risk, evidence quality, and governance constraints. The framework also incorporates decision-level explanation records and educator review, modification, rejection, and escalation mechanisms. The principal contribution is not the introduction of new standalone AI models, but the specification of how established AI capabilities can be combined within a transparent, pedagogically informed, and human-governed educational decision process.together with a staged implementation and evaluation roadmap The PCEO framework remains a conceptual and technically specified proposal; no classroom implementation or empirical effectiveness evaluation was conducted. Future work should develop and assess a working prototype through expert review, usability studies, and controlled evaluations involving educators and learners.

Keywords: AI-driven personalized learning; adaptive learning; educational decision support; explainable AI; multi-agent orchestration; human-centered AI; educational governance.

INTRODUCTION

Artificial intelligence (AI) is increasingly transforming educational systems by enabling data-driven learner analysis, adaptive instructional support, automated feedback, and personalized learning experiences. Traditional instructional approaches commonly deliver similar learning content and activities to broad groups of students despite differences in prior knowledge, learning pace, engagement, and individual support needs. AI-supported educational technologies provide an alternative by using learner data and computational models to adapt learning experiences and assist educational decision-making (Chen et al., 2020; Zawacki-Richter et al., 2019).

The development of AI-supported education has progressed from relatively static data-analysis systems toward increasingly adaptive and learner-centred environments. Early applications focused primarily on learning analytics and predictive modelling, where institutional and learner data were analysed to identify achievement patterns, predict academic performance, and detect students who might require additional support (Ouyang et al., 2022; Zawacki-Richter et al., 2019). Although these approaches strengthened evidence-based educational monitoring,

INTERNATIONAL JOURNAL OF LATEST TECHNOLOGY IN ENGINEERING, MANAGEMENT & APPLIED SCIENCE (IJLTEMAS)
ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

they were predominantly analytical and often provided limited capability for dynamically adapting instruction during the learning process.

Subsequent developments introduced AI-assisted educational applications such as intelligent tutoring systems, automated assessment, adaptive learning platforms, and natural language processing tools. Intelligent tutoring systems, in particular, have demonstrated the potential of computational techniques to provide individualized instructional support, feedback, and task sequencing (Kulik & Fletcher, 2016). More recently, generative AI and large language models have extended these capabilities by supporting conversational tutoring, automated explanation generation, learning-resource creation, and context-sensitive feedback (Kasneci et al., 2023; Tlili et al., 2023).

The broader evolution from data-driven educational technologies toward adaptive, personalized, and increasingly autonomous learning environments is summarized in Figure 1. The progression illustrates a shift from systems primarily concerned with data collection and prediction toward systems capable of learner modelling, adaptive decision-making, automated support, and overnance-aware human–AI interaction.

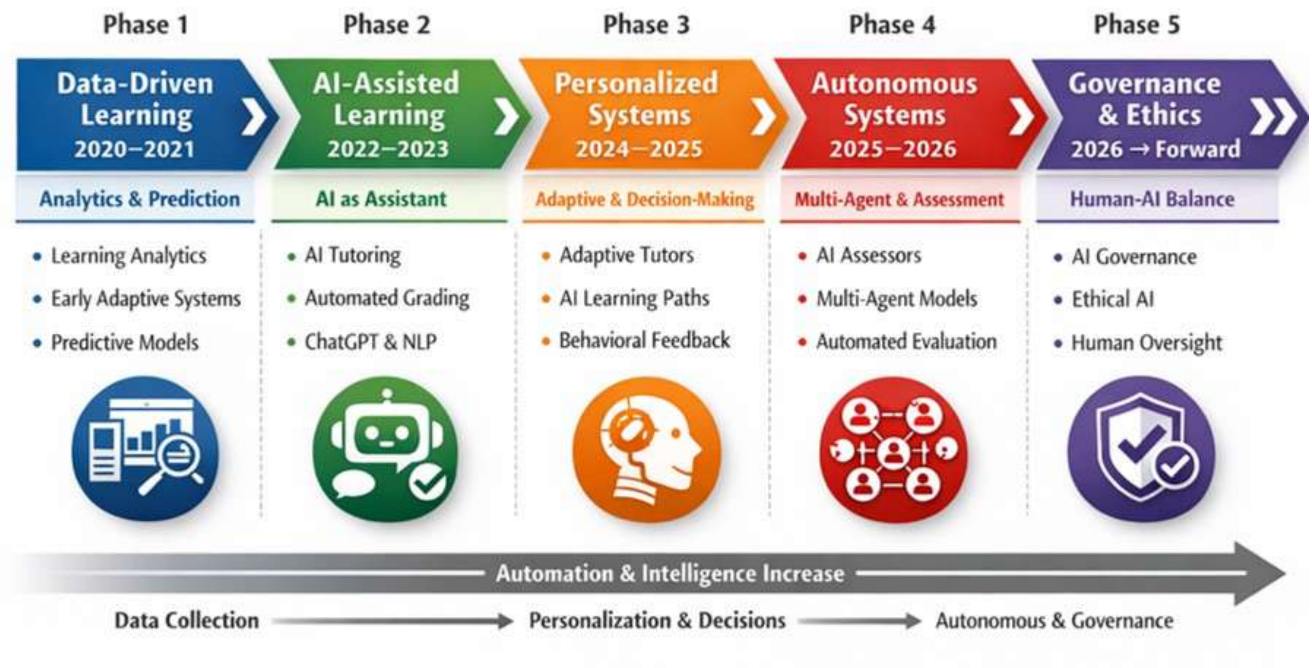

Figure 1: Evolution of learning systems from data-driven education to AI-driven personalized learning environments

As illustrated in Figure 1, increasing levels of computational intelligence have expanded the role of AI from analysing educational data to actively influencing learning pathways and instructional decisions. However, greater automation also introduces a corresponding need for transparency, pedagogical control, ethical safeguards, and meaningful human oversight. Therefore, the evolution of AI-driven education should not be understood solely as an increase in automation, but as a transition toward systems in which personalization, educational decision-making, and governance must operate together.

Contemporary personalized learning systems employ several complementary AI techniques. Learning analytics can identify patterns in learner behaviour, participation, achievement, and engagement, thereby supporting evidence-informed educational interventions (Alfredo et al., 2024). Knowledge-tracing approaches estimate how learner mastery changes across successive learning interactions and assessments, enabling more fine-grained representations of learning progress (Corbett & Anderson, 1994; Liu et al., 2022; Piech et al., 2015). Recommender and adaptive hypermedia approaches can support differentiated learning pathways by matching learners with relevant resources and activities according to their characteristics and learning state (Brusilovsky, 2001). Predictive models can identify learners at risk or estimate future academic outcomes, while adaptive

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

policy approaches can support sequential decisions regarding learning activities and interventions (Chen et al., 2020; Ouyang et al., 2022).

Generative AI has further expanded the potential scope of personalized learning by enabling interactive explanations, conversational support, automated content generation, and immediate feedback. These capabilities may increase the responsiveness of digital learning environments and provide learners with support that is more sensitive to their immediate needs (Kasneci et al., 2023; Tlili et al., 2023).

Nevertheless, generative capabilities also introduce substantial risks, including inaccurate information, hallucinated content, over-reliance, insufficient source transparency, and inappropriate automated guidance. Consequently, generative AI should be integrated into educational systems within clearly defined pedagogical, technical, and governance constraints rather than treated as an independent instructional authority (Miao & Holmes, 2023).

Despite substantial technological progress, several limitations remain in existing AI-driven personalized learning systems. First, many studies and practical systems focus on individual AI techniques—such as prediction, recommendation, knowledge tracing, or generative tutoring—without providing an explicit mechanism for coordinating their outputs within a unified decision architecture. This fragmentation can make it difficult to determine how multiple models should interact, how contradictory recommendations should be resolved, and which component should ultimately determine an educational action (Alfredo et al., 2024; Williamson & Eynon, 2020).

Second, technical performance is not equivalent to pedagogical effectiveness. A model may achieve high predictive accuracy while recommending an action that is unsuitable for the learner's current level of mastery, cognitive load, learning objectives, or instructional context. Research on AI in education has therefore emphasized the need to move beyond technology-centred optimization toward systems that integrate educational principles, learner needs, and human decision-making (Ouyang et al., 2022; Zawacki-Richter et al., 2019).

Third, explainability remains a significant challenge. Educational AI systems often produce predictions or recommendations without providing sufficiently interpretable reasons for their decisions. In educational contexts, this limitation is partic-ularly important because learners and educators may need to understand not only what the system recommends but also why the recommendation was generated, what evidence contributed to it, how confident the system is, and whether alternative actions are available.

Explainable AI in education should therefore support actionable human understanding rather than merely exposing technical model features (Khosravi et al., 2022). A related limitation concerns the relationship between AI personalization and educator decision-making. Al-though many systems focus on adapting learning experiences for individual students, educator involvement is often treated as secondary. This creates a potential disconnect in which learners receive automated recommen-dations while educators have insufficient information, authority, or mechanisms to inspect and modify those decisions.

Human-centred educational AI instead requires complementary roles in which AI supports educator judgement rather than replacing it (Alfredo et al., 2024; Holstein et al., 2019). AI-driven personalization also raises broader ethical and governance concerns. The extensive collection and analysis of learner data can cre-ate risks relating to privacy, surveillance, consent, data ownership, algorithmic bias, and unequal educational opportunities.

Learners may also become overly dependent on algorithmically selected pathways, potentially reducing autonomy or limiting exposure to diverse learning experiences. Responsible implementation therefore requires explicit safeguards for fairness, privacy, transparency, contestability, accountability, and human oversight (Foltýnek et al., 2023; Ifenthaler & Schumacher, 2016; UNESCO, 2021).

The major strengths and limitations of AI-driven personalized learning from learner and educator perspectives are summarized in Table 1.

Table 1: Strengths and limitations of AI in education from different perspectives

Perspective	Strengths	Limitations
Student	Adaptive learning, personalized content, immediate feedback, and continuous learning support	Over-personalization, dependency, limited transparency, and privacy concerns
Educator	Data-driven insights, learner monitoring, performance prediction, and targeted intervention support	Black-box decisions, limited interpretability, insufficient human control, and usability and trust concerns

As summarized in Table 1, AI-driven systems offer considerable potential for personalized instructional support and data-informed educational decision-making. However, the benefits of personalization may be constrained when systems lack pedagogical grounding, explainability, educator involvement, or enforceable governance mechanisms. These limitations indicate that effective personalized learning cannot be achieved simply by combining additional AI models or increasing the level of automation.

Accordingly, there is a need for a unified framework that coordinates multiple AI techniques through a transparent and pedagogically constrained decision process. Such a framework should maintain a dynamic representation of learner state, integrate the outputs of heterogeneous AI models, evaluate candidate educational actions against pedagogical and safety requirements, provide interpretable explanations, and preserve meaningful educator authority over high-impact or uncertain decisions.

To address this need, the present study proposes a Pedagogically Constrained Explainable Orchestration (PCEO) framework for AI-driven personalized learning and educational decision support. Rather than allowing learning analytics, knowledge tracing, recommender systems, predictive models, reinforcement-learning mechanisms, and generative AI to operate as disconnected components, PCEO coordinates their outputs through a shared learner-state representation and an explicit orchestration mechanism. The framework further incorporates pedagogical constraints, evidence-quality thresholds, governance controls, explanation records, and educator override mechanisms.

The study therefore investigates how multiple AI capabilities can be formally coordinated within a unified educational architecture while preserving pedagogical validity, explainability, learner protection, and human oversight. The proposed contribution lies not in introducing new standalone AI models, but in specifying how established AI capabilities can interact, how competing recommendations can be evaluated, when automated action should be restricted, and how decisions can remain interpretable and open to educator intervention.

LITERATURE REVIEW

Artificial intelligence has become increasingly important in higher education because of its ability to support learner modelling, prediction, adaptive content delivery, automated feedback, and educational decision-making. Existing AI-driven learning systems employ techniques such as learning analytics, knowledge tracing, recommender systems, intelligent tutoring, predictive machine learning, and, more recently, generative AI. These technologies have expanded the ability of educational systems to move beyond uniform instruction toward more personalized and data-informed learning experiences (Chen et al., 2020; Ouyang et al., 2022; Zawacki-Richter et al., 2019).

Learning analytics and predictive modelling are widely used to identify patterns in learner performance, engagement, and academic risk. Such approaches allow institutions and educators to detect students who may require intervention and to make more informed decisions based on educational data (Alfredo et al., 2024; Ouyang et al., 2022). However, many predictive systems remain focused on identifying what is likely to happen rather than determining what pedagogically appropriate action should follow. Prediction alone therefore provides limited support for real-time personalized learning unless it is connected to an adaptive decision mechanism.

Learner modelling and knowledge tracing provide a more dynamic basis for personalization. Traditional knowledge-tracing methods estimate whether learners have mastered specific skills based on their previous

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

interactions, while deep-learning approaches model more complex changes in learner knowledge over time (Corbett & Anderson, 1994; Piech et al., 2015). Recent benchmarking work has further demonstrated the increasing diversity of knowledge-tracing models available for adaptive educational applications (Liu et al., 2022). Although these approaches provide useful mastery estimates, learner knowledge alone does not fully represent the learning context. Engagement, learning risk, cognitive demand, learner preference, and the reliability of available evidence may also influence whether a particular instructional action is appropriate.

Adaptive learning and recommender systems extend learner modelling by selecting resources, activities, or learning pathways according to individual learner characteristics. Adaptive hypermedia established an early foundation for modifying educational content and navigation based on learner models (Brusilovsky, 2001), while contemporary recommendation approaches use behavioural and performance data to provide more individualized learning resources. These systems can improve the relevance of learning experiences, but a high recommendation score does not necessarily indicate pedagogical suitability. A resource may match past learner behaviour while being too difficult, repetitive, insufficiently diverse, or poorly aligned with intended learning outcomes. Consequently, personalized recommendation requires pedagogical constraints in addition to algorithmic relevance. Intelligent tutoring and generative AI have further increased the level of interaction available within personalized learn-ing environments. Intelligent tutoring systems can provide adaptive feedback and individualized instructional support, and evidence indicates that they can improve learning under suitable educational conditions (Kulik & Fletcher, 2016). Generative AI extends these capabilities through conversational tutoring, personalized explanations, automated content generation, and immediate feedback (Kasneci et al., 2023; Tlili et al., 2023). However, generative systems also introduce risks including inaccurate responses, hallucinated information, learner dependency, and inappropriate automated guidance. Their educational use therefore requires grounding, content validation, uncertainty management, and human oversight rather than unrestricted generation (Miao & Holmes, 2023).

A major concern across these AI approaches is the limited integration between technological capability and pedagogical theory. Many systems emphasize predictive accuracy, recommendation quality, or automation while providing limited evidence that adaptive decisions are explicitly guided by learning principles. Educational theories offer useful foundations for addressing this limitation. Mastery learning supports progression based on demonstrated competence (Bloom, 1968), while constructivist principles emphasize building on prior knowledge with appropriate scaffolding (Vygotsky, 1978). Self-regulated learning highlights learner choice, reflection, and progress monitoring (Zimmerman, 2002), whereas cognitive load theory emphasizes avoiding unnecessary mental overload (Sweller, 1988). Connectivism further supports exposure to diverse learning resources and knowledge networks (Siemens, 2005). However, these theories are often discussed conceptually rather than translated into explicit computational rules that influence AI-generated decisions.

Explainability and educator involvement represent additional weaknesses in existing personalized learning systems. Educational AI often produces predictions or recommendations without adequately communicating why an action was selected, how confident the system is, or what alternatives were considered. Explainable AI in education therefore requires explanations that are understandable and actionable for both learners and educa-tors rather than technical model interpretation alone (Khosravi et al., 2022). Similarly, educator involvement should extend beyond receiving dashboards or performance predictions. Human-centred educational AI requires educators to inspect recommendations, understand uncertainty, and intervene when automated decisions are inappropriate or high risk (Alfredo et al., 2024; Holstein et al., 2019). The increasing use of learner data also introduces ethical and governance concerns. Personalized systems may process sensitive information relating to academic performance, behaviour, engagement, and learner characteristics. This creates risks associated with privacy, surveillance, informed consent, algorithmic bias, data ownership, and accountability (Ifenthaler & Schumacher, 2016). Responsible educational AI therefore requires fairness, transparency, privacy protection, human oversight, and mechanisms through which automated decisions can be questioned or overridden (Foltýnek et al., 2023; UNESCO, 2021). These safeguards should not be treated as external policy considerations applied only after deployment; they should influence the decision-making process itself.

Overall, the literature demonstrates that individual AI techniques provide valuable capabilities for personalized learning, but several limitations remain. Existing approaches are frequently fragmented, with learner modelling,

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

recommendation, prediction, generative AI, explainability, and educator support operating as separate functions.

Furthermore, pedagogical theory is not consistently translated into operational decision rules, model explanations do not always explain the final educational action, and governance controls are often disconnected from real-time adaptation.The resulting research gap is therefore not the absence of AI techniques, but the absence of a clearly specified mechanism for coordinating them.

Existing research provides limited guidance on how outputs from multiple AI components should be combined, how conflicting recommendations should be resolved, when an automated action should be blocked, and how educators should retain control over uncertain or high-impact decisions. To address these limitations, this study proposes the Pedagogically Constrained Explainable Orches-tration (PCEO) framework.

The framework integrates multiple AI capabilities through a shared learner-state representation and an explicit orchestration mechanism while embedding pedagogical constraints, explainability, governance controls, and educator oversight within the decision process. The contribution therefore lies in the structured coordination and control of established AI techniques rather than in introducing another isolated AI model.

Proposed Pceo Framework: Design, Architecture, And Operational Logic

Framework Design Rationale

The proposed Pedagogically Constrained Explainable Orchestration (PCEO) framework was developed to ad-dress a central limitation identified in existing AI-driven personalized learning systems: individual AI techniques are increasingly effective at performing specialized tasks, but their outputs are often implemented independently without an explicit mechanism for coordinating them into a pedagogically valid educational decision.

Knowledge-tracing models estimate mastery, predictive models identify risk, recommender systems rank learning resources, and generative AI provides interactive feedback; however, none of these capabilities alone is sufficient to deter-mine whether a particular learning action is educationally appropriate, safe, explainable, and acceptable to the educator (Alfredo et al., 2024; Khosravi et al., 2022; Ouyang et al., 2022).

The PCEO framework therefore sepa-rates specialized AI reasoning from final educational decision authority. Instead of allowing one model to control personalization, specialized functional agents contribute different evidence to a common learner-state represen-tation. Their outputs are then coordinated by the PCEO orchestration engine, which evaluates candidate actions using pedagogical criteria, model confidence, learner context, governance requirements, and human-oversight rules.

This design choice responds to three limitations in conventional approaches. First, a single-model system usually optimizes one objective, such as predicted performance or recommendation relevance, while educational decisions require consideration of several potentially conflicting factors.

Second, loosely integrated systems may produce contradictory outputs, for example, a recommender may select a difficult resource while a learner model indicates insufficient mastery. Third, technical predictions do not automatically constitute appropriate instructional actions; they must be interpreted within pedagogical and ethical constraints (Khosravi et al., 2022; Zawacki-Richter et al., 2019).

The resulting architecture is illustrated in Figure 2. The framework integrates data acquisition, learner and context modelling, specialized AI agents, learner and educator interfaces, continuous feedback, and cross-layer governance within a closed decision cycle.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

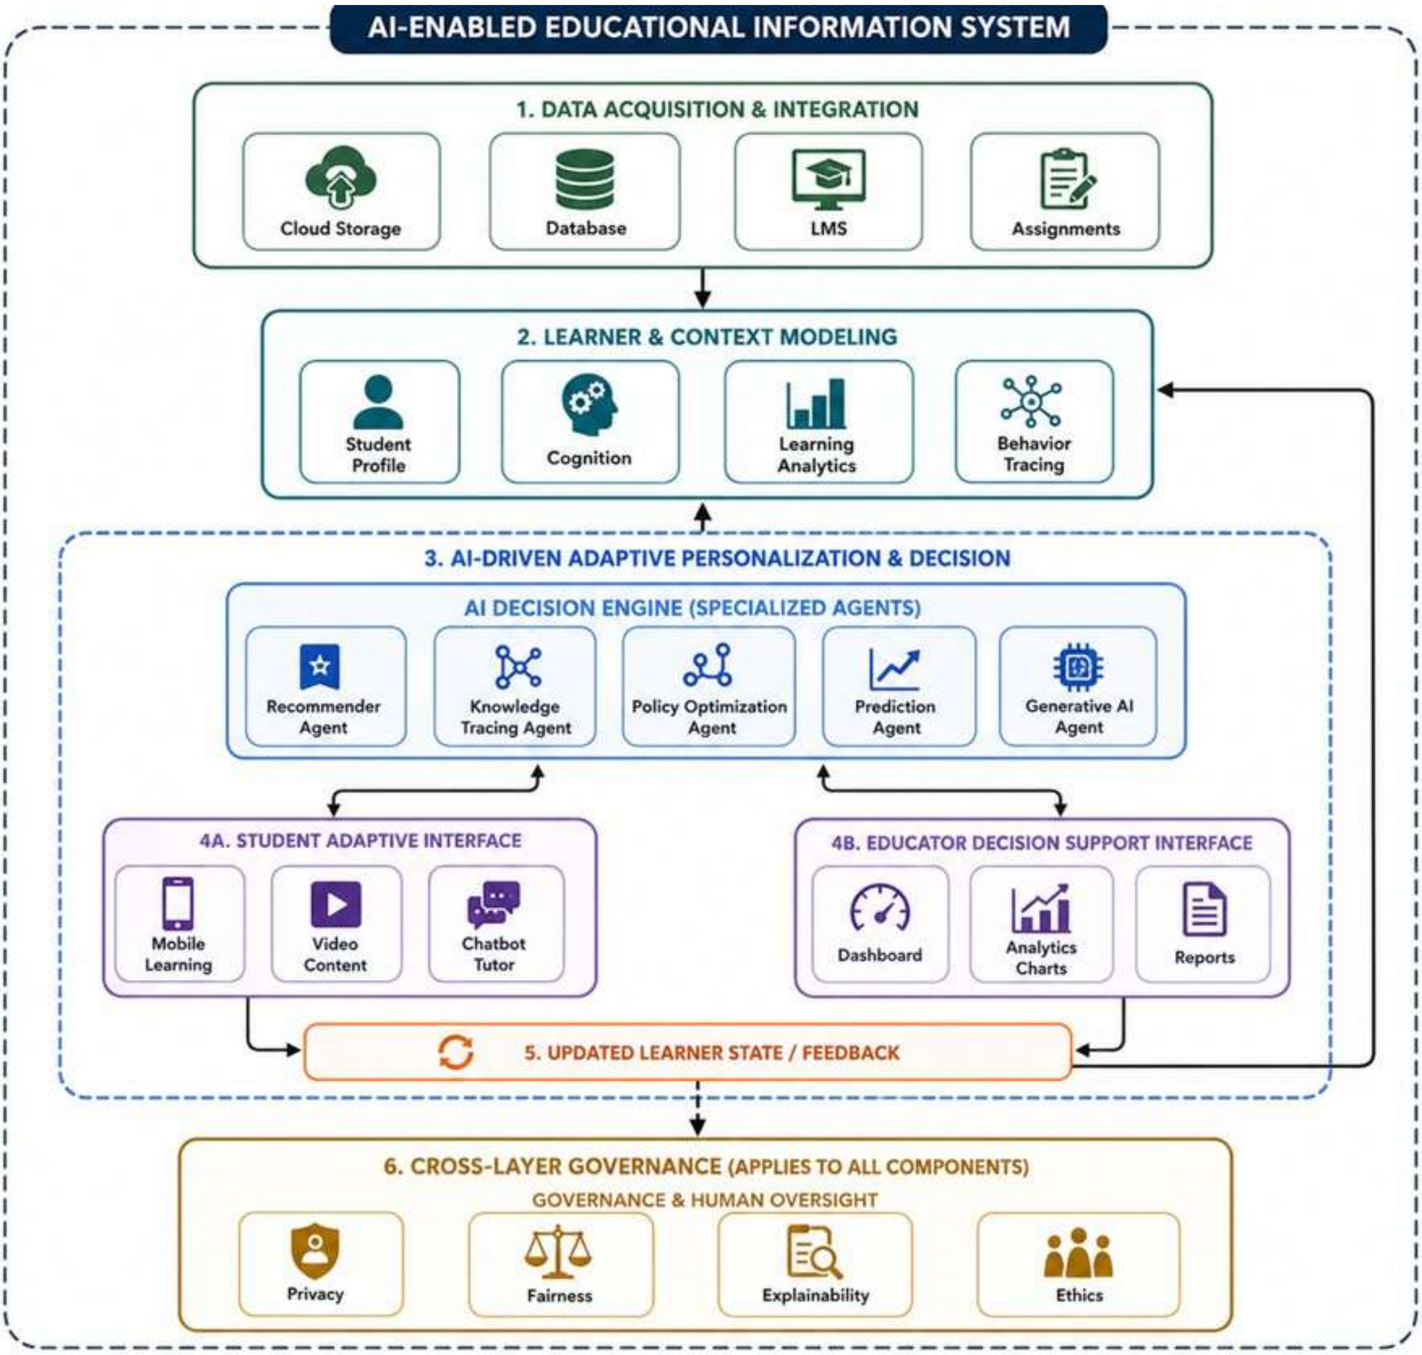

Figure 2: Architecture of the proposed Pedagogically Constrained Explainable Orchestration framework

As shown in Figure 2, the framework is not designed as a linear pipeline in which one AI model receives data and directly generates a recommendation. Instead, learner information is continuously transformed into a shared state, analysed by specialized agents, evaluated through the PCEO orchestration mechanism, and delivered through coordinated learner and educator interfaces. Feedback from subsequent learner performance and educator intervention returns to the learner model, allowing future decisions to reflect changes in learning conditions.

Data Acquisition and Shared Learner-State Representation

The framework begins with the Data Acquisition and Integration Layer, which collects information from sources such as Learning Management Systems (LMS), assessments, learner interaction logs, learning resources, institutional databases, and optional behavioural indicators. Multiple data sources are required because personalized learning cannot be represented adequately through assessment marks alone. Learner progress may also be reflected in interaction frequency, task completion, engagement patterns, previous learning activities, resource use,

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

and changes in mastery over time (Alfredo et al., 2024; Ouyang et al., 2022).

Before data are used for adaptive decision-making, the framework standardizes learner identifiers, concepts, activities, timestamps, assessment evidence, and data provenance. Missingness, recency, and reliability are also recorded.

This is important because low-quality or outdated evidence should reduce system confidence rather than silently generate a personalized action.

$$S_t = \{M_T, E_T, R_T, C_T, P_T, Q_T\}$$

where M_T represents concept-level mastery, E_T represents learner engagement, R_T represents learning k, C represents the estimated cognitive load, P_T represents the learner's preference, and Q_T represents the quality or confidence of the available evidence at interaction t.

The inclusion of these variables is deliberate. Mastery represents what the learner currently understands; engagement indicates whether the learner is actively participating; risk captures potential academic difficulty; cognitive load represents whether the proposed learning demand may exceed manageable levels; preference provides contextual information for resource delivery; and evidence quality prevents uncertain data from being treated as equally reliable. Knowledge-tracing research provides an established basis for dynamically estimating mastery rather than relying only on static historical profiles (Corbett & Anderson, 1994; Liu et al., 2022; Piech et al., 2015).

A shared learner state is important because specialized agents otherwise operate on incompatible representations. For example, a recommender may calculate content relevance while a predictive model estimates academic risk using different variables and scales. PCEO converts these heterogeneous outputs into a common decision context so that the final action reflects the learner as a whole rather than a single model prediction.

Specialized Functional Agents and Their Justification

The PCEO decision environment contains a set of specialized agents with deliberately separated responsibilities. These agents are not included simply to increase the number of AI components. Each addresses a distinct question that must be answered before a defensible personalized-learning decision can be made. The roles and justification for selecting these agents are summarized in Table 2.

Table 2: Agent Responsibilities and Standardized Outputs in the PCEO Framework

Agent	Primary Responsibility	Why It Is Required	Standardized Output
Learner Modelling Agent	Estimates concept-level knowledge and changes in mastery	Personalization requires knowing what the learner currently understands rather than relying only on previous grades	Mastery estimate (M_t) and confidence score
Engagement and Risk Agent	Detects disengagement, learning difficulties, or potential failure risks	Learners with similar mastery levels may require different interventions when their engagement or risk levels differ	Engagement estimate (E_t), risk estimate (R_t), and confidence score
Recommendation Agent	Identifies relevant and diverse learning resources or activities	Knowing a learner's current state does not determine which specific resource or activity should be presented	Relevance, suitability, and diversity scores
Policy Optimisation Agent	Evaluates the longer-term consequences and sequence of candidate actions	The locally optimal action may not produce the best learning pathway across repeated interactions	Expected action value and sequence preference

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Generative Support Agent	Produces contextual explanations, hints, feedback, and tutoring responses	Structured predictions must be translated into understandable and pedagogically appropriate instructional support	Grounded response or feedback and confidence score
Governance and Explainability Agent	Applies privacy, fairness, safety, explainability, and escalation checks	Technically suitable actions may still be unsafe, unfair, insufficiently supported, or inappropriate for automatic execution	Pass/fail gates, explanation record, and escalation status
PCEO Orchestrator	Coordinates all agent outputs and selects, approves, or escalates the final action	Separate agents may produce conflicting recommendations; therefore, an explicit mechanism is required to resolve conflicts under pedagogical constraints	Approved action, alternative action, or human-review request

As shown in Table 2, the agents are functionally complementary rather than interchangeable. The Learner Modelling Agent answers what the learner currently knows, whereas the Engagement and Risk Agent answers whether the learner is participating effectively or requires intervention. These questions cannot reliably be collapsed into a single mastery score because a learner may demonstrate moderate mastery while simultaneously showing severe disengagement.

The Recommendation Agent addresses a different problem: what learning activity or resource is suitable among available alternatives. Adaptive hypermedia and recommender approaches have long demonstrated the value of learner-sensitive resource selection (Brusilovsky, 2001). However, resource relevance alone cannot determine whether progression is pedagogically appropriate.

The Policy Optimization Agent is therefore included to consider sequential consequences. For example, immediately presenting the most challenging activity may appear optimal from a short-term recommendation perspective but may increase cognitive overload or reduce engagement across subsequent interactions. The role of this agent is not to replace pedagogical reasoning but to estimate the expected value of alternative learning sequences.

The Generative Support Agent serves yet another non-redundant function by converting structured decisions into contextual explanations, hints, feedback, or tutoring dialogue. Generative AI is particularly suited to interactive language-based support, but unrestricted generation creates risks of hallucination, inaccurate guidance, and overdependence (Kasneci et al., 2023; Miao 2& Holmes, 2023; Tlili et al., 2023). For this reason, the generative agent does not determine progression or high-impact educational decisions independently.

Finally, the Governance and Explainability Agent acts as a control mechanism rather than another predictive model. It checks whether sufficient evidence exists to justify automation and whether privacy, fairness, safety, and human-review requirements are satisfied. This separation is necessary because a model can be statistically confident while the proposed action remains ethically or pedagogically inappropriate (Foltýnek et al., 2023; Khosravi et al., 2022; UNESCO, 2021).

The PCEO Orchestrator sits above these specialized functions. Its role is critical because simply deploying multiple agents does not solve fragmentation. Without orchestration, one agent may recommend progression while another detects high risk, leaving the system without a reproducible mechanism for resolving the conflict. PCEO therefore converts multiple AI outputs into a single auditable decision process.

PCEO Orchestration and Operational Decision Logic

At each learning interaction, the specialized agents receive relevant information from the shared learner state and produce standardized outputs with confidence estimates and provenance. The PCEO orchestrator then generates a set of candidate actions such as scaffolded remediation, targeted practice, progression, enrichment, resource recommendation, generative tutoring, or educator intervention. Each candidate action a is evaluated using a multi-criteria utility function:

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

$$U(a \mid S_t) = \omega_m \Delta M + \omega_e \Delta E + \omega_p \Delta \mathrm{PedFit} + \omega_x \mathrm{Explainability} - \omega_c \mathrm{Load} - \omega_r \mathrm{Risk}$$

where ΔM represents the expected improvement in mastery, ΔE represents the expected engagement benefit, PedFit represents pedagogical suitability, Explainability represents the ability to justify the selected action, Load represents the estimated cognitive-load cost, and Risk represents the potential educational or decision-related risk. The weights ω_m, ω_e, ω_p, ω_x, ω_c, and ω_r are configurable according to the course and institutional context rather than being treated as universal constants.

The utility function addresses an important weakness in purely predictive personalization. A recommendation with a high predicted learning gain should not automatically be selected if it creates excessive cognitive load, lacks sufficient explainability, or conflicts with pedagogical requirements. Therefore, PCEO treats model predictions as evidence for decision-making rather than as the final decision itself. The candidate with the highest utility is considered only after hard constraints are evaluated:

$$a^* = \arg\max_{a} U(a \mid S_t)$$

subject to conditions such as adequate evidence quality, privacy and consent compliance, acceptable cognitive-load risk, fairness requirements, content safety, and educator approval for high-impact actions. If no candidate satisfies these conditions, the framework returns a human-review request instead of forcing an automated recommendation. The complete operational sequence is illustrated in Figure 3, beginning with learner-state construction and agent outputs, followed by candidate generation, utility evaluation, governance checks, action selection, explanation generation, dual-interface delivery, and feedback-driven state updating.

Figure 3: Operational decision flow of the PCEO orchestration engine.

As illustrated in Figure 3, governance is applied before an action is executed rather than as a post-hoc audit. This distinction is important. In conventional systems, a recommendation may first be generated and only later explained or reviewed. In PCEO, an action can be blocked before delivery when evidence is insufficient or when pedagogical, ethical, or safety constraints are violated.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Worked Illustration of the Utility Function

To make the orchestration logic concrete, consider a hypothetical, illustrative interaction in which the learner state indicates moderate mastery of a concept, $M_t \approx 0.55$, together with declining engagement and elevated cognitive load. Suppose that the orchestrator compares two candidate actions: (a) progression to the next topic and (b) scaffolded remediation with a worked example.

Progression may offer a nominally higher expected mastery gain, ΔM, because it advances curriculum coverage. However, it may also carry a higher Load penalty and lower PedFit, given the quality of the current learner evidence. In contrast, remediation may offer a smaller but more reliable ΔM, lower Load, and higher PedFit and Explainability, because it directly addresses the identified knowledge gap.

Because the utility function subtracts weighted Load and Risk terms rather than optimizing ΔM alone, remediation can obtain a higher overall utility,

$U(a \mid S_t)$,

despite producing a smaller immediate mastery gain.

This example is illustrative rather than empirically derived. The specific numerical weights,

ω_m, ω_e, ω_p, ω_x, ω_c, and ω_r,

as well as the associated decision thresholds, would need to be calibrated for a particular subject domain, learner population, and level of institutional risk tolerance during prototype development. Such calibration could be conducted through educator elicitation, historical-data analysis, or sensitivity analysis before the framework could be deployed with confidence.

Pedagogical Constraints within Agent Coordination

A central design principle of PCEO is that personalization must be constrained by educational requirements rather than optimized exclusively according to predictive performance. The framework therefore translates established pedagogical principles into operational conditions influencing action selection.

Mastery learning supports progression only when adequate evidence of competence is available (Bloom, 1968). Constructivist principles and the Zone of Proximal Development support tasks that extend prior knowledge while remaining appropriately scaffolded (Vygotsky, 1978). Self-regulated learning motivates learner choice, progress visibility, reflection, and the presentation of meaningful alternatives (Zimmerman, 2002). Cognitive load theory provides a basis for penalizing unnecessarily demanding instructional actions (Sweller, 1988), while connectivist principles motivate diversity across information sources and learning connections rather than repeatedly recommending only similar content (Siemens, 2005).

For example, an illustrative implementation may use mastery below 0.60 to trigger scaffolder remediation, mastery between 0.60 and 0.80 to support targeted practice, and mastery above 0.80 with stable engagement to permit progression or enrichment. These values are configuration examples rather than universal pedagogical thresholds and must be calibrated for the specific subject, assessment design, and learner population.

This distinction is essential. PCEO does not claim that a numerical value such as 0.80 is prescribed directly by mastery learning theory. Instead, the theory provides the educational principle progression should follow sufficient evidence of competence, while the operational threshold is empirically configured for the implementation context.

Dual Delivery, Feedback, and Human Oversight

Once an action is approved, it is delivered through coordinated learner-facing and educator-facing interfaces. The learner may receive adaptive activities, recommended resources, progress information, explanations, generative tutoring, or alternative choices. The educator receives the corresponding learner-state evidence, recommendation

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

rationale, confidence level, alternatives, and intervention options. This dual-interface design is necessary because personalization and educator decision support are frequently developed separately. Human-centred educational AI emphasizes that educators should not be reduced to passive recipients of dashboards; they require sufficient information and authority to interpret, challenge, and modify AI-supported decisions (Alfredo et al., 2024; Holstein et al., 2019). Each decision is therefore accompanied by an explanation record identifying the evidence used, reason for the recommendation, model confidence, considered alternatives, relevant provenance, and available override pathway. Educators can approve, modify, reject, or defer high-impact recommendations.

Following each learner interaction, new assessment and behavioral evidence updates the learner state from S_t to S_{t+1}. This feedback mechanism enables adaptation across repeated interactions. However, continuous adaptation is not automatically assumed to be beneficial, because erroneous predictions may propagate through the feedback loop. The framework therefore maintains confidence tracking, decision logs, safe fallback mechanisms, and human-review procedures throughout repeated adaptation.

Design Advantages over Existing AI-Based Personalized Learning Approaches

The contribution of PCEO should not be stated as "PCEO is proven better than all existing systems." Your current evidence does not support that claim. The defensible claim is that PCEO provides a more explicit and integrated design for addressing limitations that are often handled separately in existing approaches. The distinction is summarized in Table 3.

Table 3: Comparison of Common AI-Based Personalized-Learning Approaches and PCEO

Table: Design Issues in Common AI-Based Personalized Learning vs. PCEO Approach

Design Issue	Common AI-Based Personalized-Learning Approach	PCEO Approach
Learner representation	Often model-specific or based on limited variables	Shared multidimensional learner state
AI techniques	Individual or loosely connected models	Specialized agents coordinated through one orchestrator
Conflicting outputs	Often not explicitly resolved	Common utility function and constraint-based resolution
Pedagogical alignment	Frequently implicit or post-hoc	Pedagogical rules directly influence action selection
Explainability	Often explains a prediction after it is produced	Explanation attached to the complete decision pathway
Uncertainty	May be reported but not linked to action authority	Low confidence can block automation
Educator role	Dashboard, monitoring, or consultation	Explicit authority to approve, modify, reject, or escalate
Governance	Frequently treated as an external policy or post-deployment audit	Cross-layer governance gates before action execution
Adaptation	Optimizes immediate prediction or recommendation	Repeated state-aware decisions with feedback
High-impact decisions	May remain automated	Human review is required when configured thresholds are reached

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

This comparison is grounded in the design lineages that PCEO draws upon and extends rather than replaces. Adaptive hypermedia established strong foundations for learner modelling and content adaptation but gener-ally leaves conflict resolution between competing recommendations implicit (Brusilovsky, 2001). Educational explainable-AI research has developed valuable approaches to transparency and interpretability but typically explains individual model outputs rather than an end-to-end orchestrated decision (Khosravi et al., 2022). Teacher–AI orchestration research has demonstrated the value of tools that keep educators informed and in control during live instruction, but has concentrated on classroom-orchestration support rather than on coordinating multiple specialized learning-analytics and recommendation agents (Holstein et al., 2019). Human-centred learning-analytics scholarship has articulated stakeholder needs, participatory design principles, and governance considerations spanning many AI-in-education systems, but has not proposed a specific multi-agent orchestration architecture that operationalizes those principles as an executable decision process (Alfredo et al., 2024).

PCEO's distinguishing characteristic is therefore integrative rather than purely novel at the component level:

it combines agent specialization (mastery, risk, recommendation, sequential policy, generative interaction, and governance, none holding unrestricted decision authority) with pedagogy and governance participating directly in action selection rather than being added after prediction, and with a decision record that remains auditable and contestable — capturing what evidence was used, which agents contributed, why one action was selected over alternatives, how confident the system was, and whether an educator modified the outcome. These characteristics are design-level advantages rather than evidence of superior real-world effectiveness; establishing the latter requires the prototype development and empirical evaluation outlined in Section 3.9.

Implementation Roadmap and Prototype Validation Plan

Because PCEO has so far been specified only at the conceptual and architectural level, a staged implementation and validation plan is proposed to guide future work and to make explicit what would be required before effectiveness claims could be supported.

Stage 1 – Component prototyping. Each specialized agent would be implemented independently using established

techniques already validated in prior research — for example, a deep knowledge-tracing model for the Learner Modelling Agent (Piech et al., 2015), a content-based or collaborative recommender for the Recommendation Agent (Brusilovsky, 2001), and a retrieval-grounded language model for the Generative Support Agent — and evaluated against standard offline benchmarks for predictive accuracy and recommendation quality before integration.

Stage 2 – Orchestration and constraint calibration. The utility-function weights and pedagogical thresholds introduced in Sections 3.4–3.6 would be elicited from subject-matter educators (for example through structured interviews or conjoint-style preference elicitation) and refined using historical, de-identified learner-interaction data where available, with sensitivity analysis used to test how stable the orchestrator's action choices are to reasonable variation in the weights.

Stage 3 – Expert and usability evaluation. A working prototype covering a bounded subject area would be reviewed by domain educators and instructional designers for pedagogical soundness, and by learners and educators for interface usability, explanation clarity, and perceived trust, using structured protocols such as heuristic evaluation and think-aloud usability testing.

Stage 4 – Controlled and classroom evaluation. Subject to institutional ethics approval and informed consent, a controlled or quasi-experimental classroom study would compare PCEO-supported instruction against existing practice on measures including learning gains, retention, engagement, override frequency, educator workload, and explanation usefulness, with pre-registration of hypotheses and analysis plans wherever feasible.

Each stage is presented as a plan for future work rather than as evidence already collected; no component of this roadmap has yet been executed, and the results of Stages 1–4 cannot be anticipated in advance of the corresponding empirical work.

Implementation Challenges and Computational Considerations

Realizing PCEO as a working system raises practical challenges beyond the conceptual architecture. Data integration across LMS platforms, assessment systems, and institutional databases typically involves heterogeneous schemas, inconsistent identifiers, and varying data-quality standards, requiring a robust extract-transform-load pipeline and the provenance-and-recency tracking described in Section 3.2 before the shared learner state can be trusted for decision-making.

Latency is a second consideration. Because the orchestrator evaluates multiple agent outputs and governance gates before delivering a response, the end-to-end pipeline — particularly the Generative Support Agent, which may depend on a large language model — must be engineered for response times compatible with interactive use; techniques such as asynchronous pre-computation of mastery and risk estimates, caching of recommendation candidates, and lightweight governance checks on the interactive path (with heavier audits performed asynchronously) would likely be necessary in a production deployment.

Third, the framework's data requirements raise cost and infrastructure implications: maintaining per-learner state, agent confidence scores, and full decision-and-explanation logs at institutional scale requires sustained storage and processing capacity, and generative components add inference cost that scales with usage. Fourth, weight and threshold calibration (Section 3.9, Stage 2) is a non-trivial ongoing task rather than a one-time configuration step, since pedagogically appropriate weights may differ across subjects, learner populations, and institutional risk tolerance, and will need periodic review as courses and cohorts change.

Finally, governance infrastructure — access control, consent management, audit logging, and educator-override tooling — must be treated as a first-class engineering component rather than an add-on, given that Section 3.6's governance gates are intended to operate before, not after, an action is delivered. These considerations do not undermine the conceptual case for PCEO, but they indicate that a production-grade implementation would require dedicated data-engineering, MLOps, and information-security effort in addition to the modelling work described in Stage 1.

Framework Summary

The PCEO framework therefore operates as a coordinated decision architecture rather than a collection of independent AI tools. The shared learner-state representation establishes a common view of the learner; specialized agents contribute non-redundant evidence; the PCEO orchestrator resolves competing outputs; pedagogical constraints regulate instructional suitability; governance gates restrict unsafe or insufficiently justified automation; explanation records make decisions auditable; and educator override preserves meaningful human authority.

The principal contribution is therefore not the individual use of knowledge tracing, recommendation, predictive modelling, reinforcement learning, or generative AI, because these technologies already exist. The contribu-tion lies in specifying how their outputs are coordinated through a reproducible, pedagogically constrained, explainable, and human-governed decision mechanism.

METHODOLOGY

Research Design and Literature Review

This study adopted a systematic literature-informed framework development approach to investigate current applications, limitations, and design requirements of AI-driven personalized learning systems. The methodology consisted of three main stages: literature identification and screening, critical thematic synthesis, and conceptual framework development. The purpose was to derive the proposed Pedagogically Constrained Explainable Orchestration (PCEO) framework systematically from limitations and requirements identified in prior research.

The literature review was structured according to the reporting principles of PRISMA 2020, which provides guidance for transparent identification, screening, eligibility assessment, and reporting of studies in systematic

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

reviews (Page et al., 2021). Relevant literature was searched across academic databases including Scopus, Web of Science, ERIC, and IEEE Xplore, with Google Scholar used as a supplementary source for backward and forward reference searching. The search strategy combined terms related to artificial intelligence, personalized learning, adaptive learning, and higher education. The principal search expression was:

("artificial intelligence" OR "machine learning" OR "generative AI") AND ("personalized learning" OR "adaptive learning") AND ("higher education" OR university)

The primary publication period considered was 2020–2025 to capture recent developments in AI-supported education. Earlier foundational studies were retained where necessary to support established concepts such as adaptive learning, knowledge tracing, mastery learning, constructivism, self-regulated learning, cognitive load theory, and connectives.

Studies were included when they addressed AI-supported personalized or adaptive learning, were published in peer-reviewed journals or conference proceedings, and contained sufficient information about their technical approach, educational application, or system design. Duplicate records, non-academic publications, studies outside the personalized-learning scope, and studies lacking adequate methodological or conceptual relevance were excluded.

Following initial identification, records were screened based on titles and abstracts, followed by full-text assessment of potentially relevant studies. The study-selection process was documented using a PRISMA-based screening procedure.

Only actual verified search and screening counts should be reported in the final PRISMA diagram. Provisional or estimated values should not be presented as research findings.

Data Extraction and Critical Synthesis

Relevant information was extracted from the selected studies using a structured analysis approach. The review examined the principal AI techniques employed, learner-modelling approaches, adaptation mechanisms, pedagogical foundations, educator involvement, explainability mechanisms, ethical considerations, reported capabilities, and identified limitations.

The evidence was analysed using critical thematic synthesis. Rather than evaluating individual studies only according to model accuracy, the analysis focused on broader system-level issues affecting AI-driven personalized learning.

Five recurring themes were identified as particularly important to framework development:

- Fragmentation among independently operating AI techniques.
- Limited translation of pedagogical principles into adaptive decision processes.
- Insufficient explainability of AI-supported educational decisions.
- Limited educator involvement and human control.
- Concerns relating to privacy, fairness, accountability, and governance.

The synthesis also distinguished between technical capability and educational suitability. For example, a system may accurately predict learner performance or recommend highly relevant content, but this does not necessarily ensure that the resulting educational action is pedagogically appropriate, explainable, fair, or suitable for automatic execution.These recurring limitations were used as the evidence base for developing the proposed framework.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

PCEO Framework Development

The PCEO framework was developed through a requirement-driven conceptual design process. Limitations identified from the literature were translated into design requirements and subsequently mapped to corresponding architectural components and operational mechanisms.

Fragmentation among AI techniques motivated the introduction of a shared learner representation and coordinated orchestration mechanism. Instead of allowing knowledge tracing, recommendation, prediction, policy optimization, and generative AI to operate independently, the proposed framework organizes these capabilities as specialized functional components whose outputs contribute to a common decision process.

Weak pedagogical integration motivated the incorporation of explicit educational constraints. The framework draws on mastery learning, constructivist principles, self-regulated learning, cognitive load theory, and connectives to inform progression, scaffolding, learner choice, workload management, and resource diversity (Bloom, 1968; Siemens, 2005; Sweller, 1988; Vygotsky, 1978; Zimmerman, 2002).

Limited transparency in existing systems motivated the inclusion of decision-level explainability, where recommendations are intended to provide not only an output but also information about the evidence used, confidence, rationale, alternatives, and available human intervention. This design is informed by research emphasizing explainability and actionable transparency in educational AI (Khosravi et al., 2022).

Similarly, limited educator involvement motivated the inclusion of an educator decision-support mechanism that preserves the ability to review, modify, reject, or escalate AI-supported recommendations. This reflects human-centered approaches that emphasize complementary roles for educators and AI rather than unrestricted automation (Alfredo et al., 2024; Holstein et al., 2019).

Finally, privacy, fairness, ethical use, and accountability were incorporated as cross-layer governance requirements rather than being treated solely as external policy considerations (Foltýnek et al., 2023; Ifenthaler & Schumacher, 2016; UNESCO, 2021).

The PCEO framework was developed through a traceable, requirement-driven process. Evidence extracted from the reviewed studies was first synthesised to identify recurring limitations in existing AI-driven personalized learning systems, including fragmented AI components, weak pedagogical alignment, limited explainability, insufficient educator control, and inadequate governance mechanisms (Alfredo et al., 2024; Khosravi et al., 2022; Ouyang et al., 2022). As illustrated in Figure 4, each identified limitation was translated into a corresponding design requirement and then mapped to a specific component of the proposed framework. This process ensured that the framework components were supported by documented research needs rather than introduced without a clear justification.

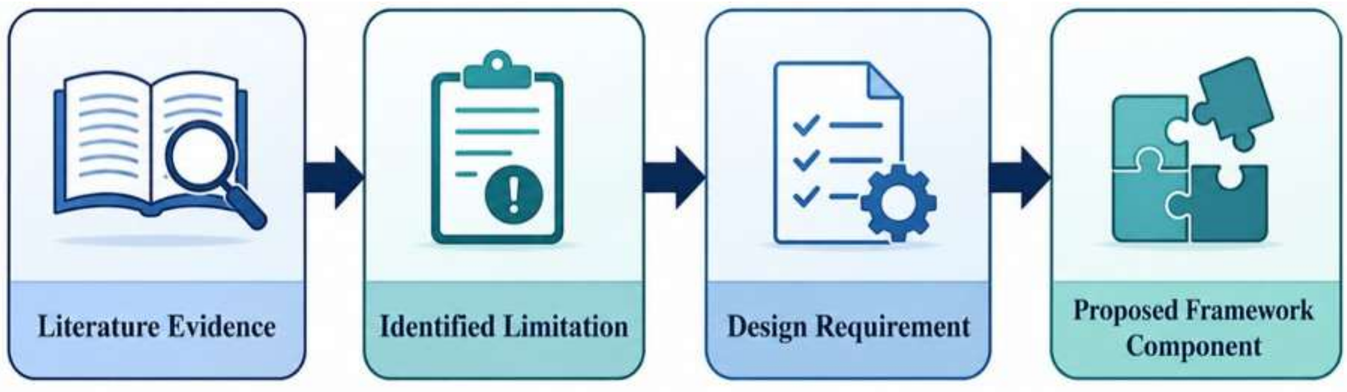

Figure 4: Requirement-driven development process of the proposed PCEO framework

As shown in Figure 4, the framework-development process maintains a direct connection between the literature evidence and the proposed architecture. For example, fragmentation among AI techniques led to the orchestration

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

component, limited pedagogical grounding led to explicit educational constraints, insufficient transparency led to explanation records, and weak human oversight led to educator review and governance controls. This traceability strengthens the conceptual justification for each component included in PCEO. The detailed architecture, specialized agents, learner-state representation, orchestration logic, pedagogical constraints, governance mechanisms, and decision flow are presented in the subsequent Proposed PCEO Framework: Design, Architecture, and Operational Logic section.

Scope of the Proposed Framework

The present study proposes and conceptually specifies the PCEO framework based on the reviewed literature. No real-world implementation, classroom experiment, learner study, or empirical effectiveness evaluation was conducted as part of this study.

Accordingly, the framework should be interpreted as a conceptual and technically specified design proposition, rather than a validated educational intervention. Claims regarding improved learning outcomes, engagement, fairness, educator workload, usability, scalability, or system performance cannot be made at this stage. Section 3.9 outlines the staged prototype-implementation and evaluation plan — expert review, prototype testing, and controlled studies involving educators and learners — required before conclusions regarding practical effectiveness could be established.

RESULTS

The literature review identified four main capabilities of AI-driven personalized learning systems: personalization, prediction, interaction, and learner monitoring. Knowledge-tracing and recommender systems support adaptive content selection, predictive models identify performance risks, intelligent tutoring and generative AI provide feedback, and learning analytics monitors learner progress (Chen et al., 2020; Ouyang et al., 2022).

Table 4: Core Capabilities of AI-Based Personalized Learning

Capability	Description
Personalization	Adapting learning resources and pathways
Prediction	Forecasting learner performance, engagement, or risk
Interaction	Providing tutoring, explanations, and feedback
Monitoring	Tracking learner behaviour and mastery

As shown in Table 5, existing systems provide valuable individual capabilities. However, these capabilities are often implemented separately rather than coordinated within a unified educational decision process. The critical synthesis identified five recurring limitations: fragmented AI components, weak pedagogical alignment, limited decision explainability, insufficient educator control, and inadequate privacy, fairness, and accountability mechanisms (Alfredo et al., 2024; Holstein et al., 2019; Khosravi et al., 2022).

These findings revealed several tensions, including predictive accuracy versus pedagogical suitability, automation versus educator judgement, personalization versus learning diversity, and data use versus learner privacy. Such tensions show that technical performance alone is insufficient for responsible personalized learning. The identified limitations were translated into PCEO design requirements, as summarized in Table 6.

Table 5: Literature Findings and Corresponding PCEO Responses

Literature Findings and PCEO Responses

Literature Finding	PCEO Response
Fragmented AI techniques	Central orchestration engine

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Narrow learner representation	Shared multidimensional learner state
Weak pedagogical alignment	Explicit pedagogical constraints
Limited transparency	Decision-level explanation card
Insufficient educator control	Review and override mechanisms
Privacy and fairness concerns	Cross-layer governance gates

As illustrated in Figure 4, the framework-development process moves from literature evidence to identified limitations, design requirements, and proposed framework components. This confirms that the PCEO components were derived from recurring research gaps rather than selected arbitrarily.

The principal result of this study is therefore a literature-informed conceptual framework that coordinates established AI techniques through pedagogical constraints, explainability, governance, and educator oversight. together with a staged plan for implementing and evaluating that framework. The findings support the conceptual

need for PCEO, but they do not demonstrate improved learning outcomes, engagement, fairness, usability, or scalability. Empirical evaluation remains future work.as outlined in Section 3.9.

DISCUSSION

The findings show that AI-driven personalized learning systems provide valuable capabilities in learner modelling, prediction, recommendation, tutoring, and progress monitoring. However, these capabilities are frequently implemented as separate technical functions rather than as parts of a coordinated educational decision process. This fragmentation creates a practical problem: a system may accurately estimate learner mastery or recommend a relevant resource without determining whether the resulting action is pedagogically appropriate, sufficiently explainable, ethically acceptable, or suitable for automatic execution (Alfredo et al., 2024; Ouyang et al., 2022).

The proposed PCEO framework addresses this issue by separating the outputs of individual AI components from the final educational decision. Knowledge tracing, prediction, recommendation, policy optimization, and generative support perform specialized functions, while the orchestration mechanism combines their outputs using a shared learner representation. This design is conceptually stronger than relying on a single model because educational decisions involve multiple and sometimes conflicting considerations. For example, a recommendation may be technically relevant but inappropriate when learner mastery is low, engagement is declining, evidence quality is weak, or cognitive demand is excessive.

As summarized in Table 5, each PCEO component responds to a specific limitation identified in the literature. Fragmented AI techniques motivated the orchestration engine, narrow learner representations motivated the shared learner state, weak pedagogical grounding motivated explicit educational constraints, and limited transparency motivated decision-level explanation records. Similarly, insufficient educator control motivated review and override functions, while concerns regarding privacy, fairness, and accountability motivated cross-layer governance mechanisms. The requirement-driven process illustrated in Figure 4 therefore provides traceability between the literature findings and the proposed architecture.

A central contribution of PCEO is the integration of pedagogical principles into the decision process. Existing educational AI research often emphasizes prediction accuracy or recommendation relevance, but technical accuracy does not guarantee educational suitability. The framework addresses this weakness by using mastery learning, constructivist scaffolding, self-regulated learning, cognitive load theory, and connectives as design constraints rather than decorative theoretical references (Bloom, 1968; Siemens, 2005; Sweller, 1988; Vygotsky, 1978; Zimmerman, 2002). This means that an adaptive action should be selected not only because a model assigns it a high score, but also because it is consistent with the learner's current competence, manageable cognitive demand, learning goals, and need for choice and resource diversity.

The framework also strengthens the role of explainability. In many existing systems, explanations are limited to individual model predictions. However, explaining why a model classified a learner as being at risk is not the

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

same as explaining why the system selected a particular intervention. PCEO therefore treats explainability as part of the complete decision record, including the evidence used, confidence level, rationale, alternatives, and available human intervention. This approach is consistent with educational explainable AI research, which emphasizes understandable and actionable explanations for learners and educators (Khosravi et al., 2022).

Educator involvement is another important design consideration. Existing systems frequently provide dashboards or analytics but give educators limited authority over automated decisions. PCEO instead adopts a human-centered approach in which educators can review, modify, reject, or escalate recommendations. This supports teacher–AI complementarity by positioning AI as a decision-support mechanism rather than an independent instructional authority (Alfredo et al., 2024; Holstein et al., 2019). Such oversight is particularly important for uncertain, high-impact, or context-sensitive decisions where the educator may possess information that is unavailable to the system.

Governance is also embedded throughout the framework rather than added after a decision has been made. Privacy, consent, fairness, accountability, evidence quality, and uncertainty can influence whether an action is approved, blocked, or escalated. This is important because increasingly detailed learner modelling may improve personalization while simultaneously increasing risks related to surveillance, data ownership, bias, and learner autonomy (Foltýnek et al., 2023; Ifenthaler & Schumacher, 2016; UNESCO, 2021). PCEO therefore treats governance as an operational requirement rather than a separate ethical checklist.

The proposed framework should not, however, be interpreted as empirically superior to existing personalized learning systems. The present study developed PCEO through literature synthesis and conceptual design only. No classroom implementation, learner experiment, educator study, or technical performance evaluation was conducted. Consequently, the study does not demonstrate improved learning outcomes, engagement, trust, fairness, usability, educator workload, or scalability.

The main contribution is therefore conceptual and architectural. PCEO provides a clearer specification of how established AI capabilities may be coordinated through a shared learner state, pedagogical constraints, explanation mechanisms, governance controls, and educator oversight.together with the staged implementation roadmap and computational considerations set out in Sections 3.9 and 3.10. Future work should implement a prototype following that roadmap and and evaluate it through expert review, usability testing, and controlled classroom studies. Such evaluation should examine learning gains, retention, learner agency, explanation usefulness, educator workload, fairness, override frequency, and system reliability before practical effectiveness can be claimed.

CONCLUSION

This study examined the current landscape of AI-driven personalized learning and identified recurring limitations in the integration, pedagogical grounding, explainability, educator involvement, and governance of existing systems. Although technologies such as learning analytics, knowledge tracing, recommender systems, predictive models, intelligent tutoring systems, and generative AI provide valuable capabilities, they are often implemented as isolated components rather than as part of a coordinated educational decision process. To address these limitations, the study proposed the Pedagogically Constrained Explainable Orchestration (PCEO) framework. The framework coordinates specialized AI components through a shared learner-state representation, an explicit orchestration mechanism, pedagogical constraints, decision-level explanations, governance controls, and educator review and override. Its main contribution lies not in introducing new standalone AI techniques, but in specifying how established capabilities can be combined within a transparent, pedagogically informed, and human-governed architecture.

The literature synthesis supports the conceptual need for such a framework, particularly where technical accuracy alone is insufficient to justify educational action. PCEO therefore positions AI as a decision-support mechanism rather than an autonomous instructional authority, while maintaining learner protection, explainability, and educator responsibility.

However, the framework remains a conceptual and technically specified proposal. It has not yet been implemented or empirically evaluated in real educational settings. Future research should therefore develop a working prototype and assess it through expert review, educator usability studies, learner trials, and controlled classroom evaluation. Such studies should examine learning outcomes, retention, engagement, fairness, explanation usefulness, educator workload, system reliability, and scalability before claims of practical effectiveness can be made.

Conflict Of Interest

The authors declare that they have no known professional, or personal conflicts of interest that could have influenced the research reported in this manuscript.

Author Contributions

First Author: Conceptualization, methodology, investigation, data analysis, and original manuscript preparation. Second Author: Software development, validation, data visualization, and manuscript review.

Third Author: Supervision, project administration, validation, and manuscript editing. All authors read and approved the final version of the manuscript.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the individuals, institutions, and organizations that provided support, resources, data, technical assistance, or feedback during this research.

REFERENCES

1. Alfredo, R., Echeverria, V., Jin, Y., Yan, L., Swiecki, Z., Gašević, D., & Martinez-Maldonado, R. (2024). Human-centred learning analytics and AI in education: A systematic literature review. Computers and Education: Artificial Intelligence, 6, 100215.
2. Bloom, B. S. (1968). Learning for mastery. Evaluation Comment, 1(2), 1–12.
3. Brusilovsky, P. (2001). Adaptive hypermedia. User Modeling and User-Adapted Interaction, 11, 87–110.
4. Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. IEEE Access, 8, 75264–75278.
5. Corbett, A. T., & Anderson, J. R. (1994). Knowledge tracing: Modeling the acquisition of procedural knowledge.
6. User Modeling and User-Adapted Interaction, 4, 253–278.
7. Foltýnek, T., Bjelobaba, S., Glendinning, I., Khan, Z. R., Santos, R., Pavletic, P., & Kravjar, J. (2023). ENAI recommendations on the ethical use of artificial intelligence in education. International Journal for Educational Integrity, 19, Article 12.
8. Holstein, K., McLaren, B. M., & Aleven, V. (2019). Co-designing a real-time classroom orchestration tool to support teacher–AI complementarity. Journal of Learning Analytics, 6(2), 27–52.
9. Ifenthaler, D., & Schumacher, C. (2016). Student perceptions of privacy principles for learning analytics.
10. Educational Technology Research and Development, 64(5), 923–938.
11. Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274.
12. Khosravi, H., Buckingham Shum, S., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., & Gašević, D. (2022). Explainable artificial intelligence in education. Computers and Education: Artificial Intelligence, 3, 100074.
13. Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review.Review of Educational Research, 86(1), 42–78.
14. Liu, Z., Liu, Q., Chen, J., Huang, S., Tang, J., & Luo, W. (2022). pyKT: A Python library to benchmark deep learning-based knowledge tracing models. Advances in Neural Information Processing Systems, 35.
15. Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

16. Ouyang, F., Zheng, L., & Jiao, P. (2022). Artificial intelligence in online higher education: A systematic review of empirical research from 2011 to 2020. Education and Information Technologies, 27(6), 7893–7925.
17. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71.
18. Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. Advances in Neural Information Processing Systems, 28, 505–513.
19. Siemens, G. (2005). Connectivism: A learning theory for the digital age. International Journal of Instructional Technology and Distance Learning, 2(1), 3–10.
20. Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12(2), 257–285.
21. 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, Article 15.
22. UNESCO. (2021). Recommendation on the ethics of artificial intelligence.
23. Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Harvard University Press.
24. Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education.
25. Learning, Media and Technology, 45(3), 223–235.
26. Zawacki-Richter, O., Mar´ın, V. I., Bond, M., & Gouverneur, F. (2019). A systematic review of research on artificial intelligence applications in higher education—Where are the educators? International Journal of Educational Technology in Higher Education, 16, Article 39.
27. Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. Theory Into Practice, 41(2), 64–70.