

Enhanced Explainable Multi-Agent Generative Artificial Intelligence Framework for Personalized Education

Hesham M. A. Abdullah¹, Elham Alzain^{2*}

Al-Fayha Private College, Jubail, Saudi Arabia¹

Applied College, King Faisal University, Al-Ahsa, 31982, Saudi Arabia²

Abstract—Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs) have enabled personalized education via intelligent tutoring, adaptive learning, automated assessments, and personalized feedback. Nevertheless, educational AI systems are still confronted with challenges, including high computational complexity, long inference time, unreliable content generation, limited explainability, and ineffective multi-agent orchestration. To this end, this study proposes LITE-Trust, a Lightweight Trust-Aware Explainable Multi-Agent Generative Artificial Intelligence Framework for personalized education, which consists of learner profiling, Knowledge State Estimation (KSE), the Educational Library and Readers Optimization (ELRA) algorithm, relevance-based selective agent orchestration, trust evaluation, and an Explainability Engine. The ELRA algorithm is applied to optimize learner-specific educational configurations by taking learning requirements, trust, and computational efficiency into account, and selective orchestration is conducted to reduce computational overhead by only activating relevant agents. Experiments on the EdNet and Open University Learning Analytics Dataset (OULAD) demonstrate that LITE-Trust achieves 10.5/9.5ms latency, 80.5/81.3 MB memory, and 1.92 M parameters, which is computationally more efficient than the compared methods. In addition, LITE-Trust attains 95.0%/94.6% fidelity, 94.1%/93.4% explanation quality, and 92.6%/92.1% interpretability on the two datasets. Besides, the ELRA algorithm significantly outperforms existing algorithms in terms of various metrics. Overall, the experimental results verify that LITE-Trust can effectively balance personalization, learning quality, trust, explainability, and computational efficiency, thereby facilitating the practical application of lightweight and trustworthy GenAI systems in personalized education.

Keywords—Generative artificial intelligence; artificial intelligence in education; personalized learning; multi-agent systems; explainable artificial intelligence; trustworthy AI; educational optimization; lightweight artificial intelligence; library and readers algorithm; learning analytics

I. INTRODUCTION

Artificial Intelligence in Education (AIED) has evolved from rule-based educational systems into Intelligent Tutoring Systems (ITSs), personalized learning, Generative AI, and LLMs [1, 2]. These advances allow the creation of content, recommendations, and feedback and offer tutoring services to learners individually or in groups. Still, the growing complexity of AIED and the use of large-scale models and

many agents have made education-related artificial intelligence challenging to develop and control.

Current educational AI systems suffer from several limitations. Large-scale transformer-based models increase the computational complexity, memory consumption, inference time, and cost of deploying these systems. Multi-agent educational systems execute multiple agents without considering relevance to the current educational context, thus introducing unnecessary costs [3]. In addition, educational tutoring systems currently lack reliable mechanisms to assess trust and hallucination, potentially propagating unsubstantiated or factually incorrect information to the learner. Finally, these systems have limited explainability, as neither the learners nor instructors are often privy to the decision-making process of the AI. Most importantly, existing optimization methods are typically single-point, only considering one optimization criterion (e.g., prediction accuracy or learning path optimization) rather than jointly optimizing prediction, learning path adaptation, trust, explainability, agent coordination, and computational cost.

These limitations motivate this work by highlighting several research challenges, including single-agent, static multi-agent, limited trust-aware optimization, insufficient explainability, high computational cost, the absence of lightweight educational AI, and limited educational optimization. Table I explains the comparison of proposed and existing approaches.

TABLE I. COMPARISON OF EXISTING APPROACHES AND LITE-TRUST

Methods	Personalization	Multi-Agent	Trust	XAI	Lightweight	Optimization
Existing method	✓	Partial	✗	Partial	✗	Partial
LITE-Trust	✓	✓	✓	✓	✓	✓

The proposed LITE-Trust is a lightweight multi-agent educational AI framework that comprises 1) learner profiling, 2) KSE, 3) ELRA, 4) dynamic agent orchestration, 5) trust evaluation, and 6) explainability. The ELRA optimizes learner-specific educational configurations while relevance-based pruning minimizes unnecessary agent execution. On the other hand, trust evaluation enhances the credibility of produced

*Corresponding author.

educational content, and the explainability engine offers learner-specific rationales. The major contributions include the following:

- A lightweight multi-agent educational AI framework that incorporates personalization, optimization, trust, and explainability features.
- The ELRA algorithm for personalized education optimization.
- A trust-aware evaluation method for enhancing the reliability of generated educational content.
- An explainability engine that enables transparent personalization decisions.
- A dynamic agent orchestration technique that minimizes computational overhead while maximizing the accuracy of personalization.
- A comprehensive evaluation of the performance in terms of learning, trust, explainability, optimization, and computational efficiency.

The rest of this study is structured as follows: Section II describes the related work. Section III presents the LITE-Trust framework. Subsequently, Section IV provides an experimental analysis and discussion. Finally, Section V concludes this study and identifies limitations and future directions.

II. RELATED WORKS

A. AI and Generative AI in Education

AI and Generative AI methods efficiently enhance collaborative learning, AI-enabled teaching, and lesson preparation in educational systems. The following presents a survey of AI and Generative AI in educational systems. Alyoussef et al. [4] and Algerafi et al. [5] present a technology acceptance model-based AI adoption architecture for cooperative learning. The system incorporates perceived accessibility, recognized usefulness, engagement factors, and trust attitude to evaluate AI adoption. To evaluate the relationships between these factors, this the study uses structural equation modeling. The system enhances students' behavioural intention and AI-supported technologies in learning environments. Oh et al. (2025) present the AI-based MATH41 system for mathematics teacher education. This approach is employed for generating teaching materials for students and teaching with the aid of Generative AI and MATH41 methods. Here, the learning materials are prepared by the system, which helps to enhance lesson planning skills and AI-enabled learning.

Overall, the discussed research highlights the potential of using AI and GenAI for enhancing engagement, adoption, and content development; however, they are focused on different aspects of using these technologies. The adoption-centric studies explore behavioural and psychological aspects, while GenAI research mainly concentrates on content and teaching materials. Hence, there is a gap in the studies, which implies the need for further investigations. Specifically, there is a need for studies exploring ways of combining learner state

adaptation and using GenAI for generating decisions with high computational efficiency and reliability.

B. Personalized and Adaptive Learning

Personalized and dynamic learning methods utilized AI to analyze learners' knowledge, ability, and learning process to provide customized learning content, dynamic assessments, and personalized learning pathways. These methods effectively improve learner participation, learning efficacy, and academic excellence. The following section presents a survey of personalized and responsive learning works. Pradeesh et al. (2025) present an AI-supported personalized learning method by employing a knowledge-centric concept assessment. The subsystem incorporates retrieval-augmented generation, exercise-aware knowledge tracing, and Generative AI to analyze students' knowledge deficiencies, generate responsive evaluations, and provide personalized recommendations. Therefore, this research enhances learning outputs using concept-level assessments. Ma et al. [8] introduce an AI-based adaptive language learning system that consolidates natural language processing and a recurrent neural network to generate personalized learning approaches. The component analyses students' linguistic proficiency, speech patterns, and grammar to provide dynamic learning content, deliver real-time feedback, and recommend adaptive learning pathways. Thus, this research enhances student participation and English proficiency. Lou et al. [9] present an adaptive educational path-planning methodology that implements reinforcement learning and a cognitive graph to produce personalized learning paths. The subsystem continuously monitors learners' cognitive states, evaluates learning performance, and enhances content sequencing using predictive analytics. Thus, the presented system improves learning efficiency and educational outputs.

These approaches realize personalization based on knowledge tracing, learner proficiency analysis, and reinforcement-optimized path selection. Nevertheless, most of the proposed approaches concentrate on content adaptation and path personalization, while trustworthiness, explainability, and computational efficiency are insufficiently considered, which emphasizes the necessity of the developed approach, which is based on learner knowledge trace personalization.

C. Multi-Agent Systems in Education

Multi-Agent Systems (MAS) enable academic applications by allocating learning, tutoring, recommending, and assessment functions among several agents. Viswanathan et al. (2022) introduce a multi-agent learning system that coordinates multiple agents of personalized educational activities. Chang et al. [11] suggested a multi-agent approach to personalize education by dividing educational tasks among several agents. Rida et al. (2025) attempted to perform a multi-agent approach to achieving personalized educational services by assigning educational tasks to cooperating agents. Overall, the authors demonstrated that multi-agent approaches could be applied to perform specialized tasks, specifically in education. However, the majority of the works above employed a multi-agent approach without jointly considering trust evaluation, explainability, computational simplicity, and optimization, which limits adaptive and personalized learning by requiring extra agents to perform learning tasks, thus imposing

substantial overhead. As such, there is a need to develop a lightweight multi-agent architecture that targets accomplishing educational tasks while incorporating trust evaluation and explainability to ensure high-quality personalized learning.

D. Trustworthy and Explainable AI in Education

Trustworthy and explainable AI is necessary in education to confirm that AI-enabled decisions are transparent, responsible, explainable, and dependable for students and teachers. Latest studies focus on developing trustworthy governance, interpretable academic performance prediction, and reliable knowledge tracking strategies. The following is a survey of trustworthy and explainable AI in education. Evangelista et al. [13] introduce an innovative approach to a policy-as-code system that ensures trustworthy AI management in higher education. The component tracks institutional requirements and converts them into computable policies and uses a gatekeeper to evaluate system performance, calibration, fairness, explainability, and system limits. It additionally integrates fairness mitigation, SHapley Additive exPlanations (SHAP)-based interpretability, uncertainty-aware decisions, and tamper-evident audit procedures. Hence, the system allows explainable, accountable, and implementable AI governance during the development and deployment process. Raji et al. [14] introduce an explainable Machine Learning (ML) method to forecast the Deaf and Hard of Hearing learners' performance. The study deployed ML methods with Local Interpretable Model-agnostic Explanations (LIME) and SHAP to explain the forecasting results and detect important academic performance features. The combination of eXtreme Gradient Boosting (XGBoost), ExtraTrees, and Random Forest models resulted in 92.99% accuracy. Lu et al. (2024) proposed a trustworthy knowledge tracking model for an AI-enabled tutor system. Then, proposed three-level solutions to improve transparency, accountability, and interpretability through the integration of AI explanation techniques with educational systems. In turn, technical explanations are transformed into natural learning outcomes and presented through the ITS interface. Experimental evaluation shows that the presented model enhanced students' and teachers' trust and knowledge of the system.

These works relate to trustworthiness at various levels, from governance to prediction explainability and knowledge tracking. Yet, the emphasis is on providing an appropriate explanation or validation of each AI-based decision. These studies do not consider the combined evaluation of trust in the system with optimization of individual decisions and coordinated adaptation of agents. Consequently, it appears that the development of an integrated approach for a trustworthy, explainable, and learner-based system is needed.

E. Lightweight AI for Educational Applications

Lightweight AI models have emerged as capable approaches for educational applications, which enable real-time learner assessment while requiring lower computational resources. The following is a survey of lightweight AI for educational applications. Radharani et al. [16] present a light attention-based framework for detecting student attentiveness based on their facial expressions. MobileNet, together with channel and spatial attention, is deployed to identify facial

features such as eye closure, mouth movements, and gaze. It achieves 91.2% accuracy with fewer parameters, which supports real-time classroom deployment [17]. Yu et al. (2025) present an IoT and edge-assisted system for computer-assisted learning. Here, IoT sensors are employed to record the engagement level of the students, while the edge devices process the data locally to provide immediate feedback with the aid of a lightweight stacked Long Short-Term Memory (LSTM) model. The system provides students with real-time feedback with 99.3% accuracy. Qorich et al. [18] introduce a context-based AI-enabled question generation approach for academic assessment that employs a lightweight transformer model. The subsystem combines a Convolutional Neural Network (CNN) module with a Bidirectional and Auto-Regressive Transformer (BART) base to enhance context-aware representation while preserving lower computational complexity; hence, this framework is suitable for a real-time educational environment.

These studies reduce computational costs by utilizing edge processing and lightweight transformers while ensuring effective educational outcomes. However, the efficiency is demonstrated in individual components, while the system-level factors, including agent execution, optimization, and trustworthiness, are limited [19]. Thus, more attention should be dedicated to developing educational frameworks that guarantee efficient model and end-to-end systems.

F. Metaheuristic Optimization for AI-Based Education

Optimization strategies employed in AI-powered learning include the use of heuristic search and evolutionary algorithms in optimizing learning pathways, content sequencing, learner modelling, and academic decision support. Most of the existing works used optimization algorithms like Ant Colony Optimization (ACO), Genetic Algorithm (GA), Particle Swarm Optimization (PSO), and Library and Readers Algorithm (LRA) [22]. These algorithms enhance optimize learning approaches, educational assessment, quality of recommendations, and adaptive educational processes by identifying optimal solutions under multiple learning objectives. Subsequently, they enhance personalization, learner participation, and overall educational performance. However, most previous optimization methods mainly focus on optimizing individual educational tasks while needing iterative optimization and parameter tuning, which improves computational overhead and execution time. Furthermore, they provide limited support for lightweight multi-agent coordination, trust-based education content assessment, interpretable AI recommendations, dynamic prompt optimization, and computationally effective personalized learning [23]. The abovementioned limitations emphasize the requirement for an integrated optimization system capable of balancing personalization, trustworthiness, transparency, and computational efficiency in AI-driven educational technologies.

G. Research Gap and Positioning of the Proposed Framework

The literature review demonstrates that previous AI-enabled personalized learning systems have intelligent learner profiling, advanced tutoring, content recommendation, and

learning path generation. However, they do not adequately address high computational overhead, ineffective trust evaluation, static multi-agent coordination, limited explainable recommendations, and optimization methods that mainly focus on individual educational tasks instead of an integrated learning system [24]. Furthermore, recent methods offer limited optimization of AI agent selection, learning path adaptation, prompt generation, trust analysis, and computational resource allocation simultaneously, which reduces their practical deployment and scalability in resource-limited educational environments.

To address these challenges, the presented LITE-Trust system designs a lightweight trust-aware explainable multi-agent generative AI architecture optimized using the presented ELRA algorithm. In contrast to existing approaches, ELRA automatically coordinates educational agents, executes trust-based optimization, adapts personalized learning pathways, optimizes prompt selection, and minimizes computational complexity while integrating explainability and academic integrity assessment. Subsequently, the presented system positions itself as a unified, lightweight, trustworthy, and computationally effective approach for advanced AI-assisted personalized learning.

For the purpose of comparison, the capabilities in Table II are classified according to their explicit implementation in the reviewed studies. A ✓ indicates that the capability is directly implemented as a core component, X indicates that it is not obviously addressed, and Partial indicates that the capability is considered indirectly or to a limited extent. The classification is based on the methods and objectives reported in the corresponding studies instead of inferred from their application context. Table II indicates that most of the studied works address some specific capabilities rather than the combination of their implementation. Personalization and optimization are the most commonly performed capabilities. Some papers address the coordination of multi-agent systems, but it is not fully explored, and only a few works mention multi-agent coordination, trust evaluation, XAI, and lightweight processing. In contrast, the LITE-Trust system combines several capabilities, which are partially addressed in the existing solutions, namely KSE-based personalization, ELRA optimization, dynamic multi-agent coordination, trust evaluation, XAI, and lightweight. However, it should not be assumed that other systems do not possess the mentioned capabilities.

TABLE II. RESEARCH GAP AND POSITIONING OF LITE-TRUST

Approach	Personalized	Multi-Agent	Trust	XAI	Lightweight	Optimization
Alyoussef et al. (2025)	✓	X	Partial	X	X	X
Algerafi et al. (2023)	✓	X	X	X	X	X
Oh et al. (2025)	✓	X	X	X	X	X
Pradeesh et al. (2025)	✓	X	Partial	X	X	X
Ma et al. (2025)	✓	X	X	X	X	X
Lou et al. (2025)	✓	X	X	X	X	✓
Viswanathan et al. (2022)	✓	✓	X	X	X	X
Chang et al. (2026)	✓	✓	X	X	X	X
Rida et al. (2025)	✓	✓	X	X	X	X
Evangelista et al. (2026)	X	X	✓	✓	X	Partial
Raji et al. (2024)	X	X	Partial	✓	X	X
Radharani et al. (2026)	X	X	X	X	✓	✓
Yu et al. (2025)	✓	X	X	X	✓	✓
Kalwar et al. (2025)	✓	X	X	X	X	✓
Imamah et al. (2024)	✓	X	X	X	X	✓
[21] Zeng et al. (2022)	✓	X	X	X	X	✓
Proposed LITE-Trust	✓	✓	✓	✓	✓	✓

III. PROPOSED LITE-TRUST FRAMEWORK

A. Overall Architecture

The presented LITE-Trust framework initially collects learner information, including demographic characteristics, academic history, prior knowledge, learning preferences, assessment performance, engagement patterns, and interaction logs, to construct a dynamic learner profile as shown in Fig. 1.

The KSE subsystem then analyzes the obtained profile to estimate learners' competency levels and identify learning gaps. These outputs are then processed by the proposed ELRA, which identifies candidate educational policies, evaluates trust-aware recommendations, and optimizes educational decisions through personalized enhancement. Based on the optimized strategy, the lightweight multi-agent collaboration module adaptively activates and collaborates with the most suitable educational agents, which include the tutor, assessment,

feedback, recommendation, and academic integrity agents, based on the current learning context. Then, the explainability component analyzes the produced recommendations to provide transparent educational justifications, while the academic integrity module analyzes responsible AI usage, originality, critical thinking, and authentic learning behaviour. Finally,

personalized learning recommendations are offered to learners, and continuous learner interactions are used to update learner profiles and knowledge state estimation, which enables adaptive, trustworthy, explainable, and computationally effective AI-enabled personalized learning.




Fig. 1. Overall architecture of the LITE-Trust framework.

B. Student Profile Construction Modules

The student profile construction subsystem establishes a dynamic and comprehensive learner profile, supporting personalized educational decision-making within the LITE-Trust system. Previous AI-enabled learning systems often depend on static learner profiles or limited educational indicators, which makes it difficult to extract learners' developing knowledge, engagement, behavioural patterns, and individual learning preferences, which frequently leads to ineffective personalisation. To resolve the existing

shortcomings, the proposed module combines heterogeneous learner data, which includes demographic characteristics (D_i), prior knowledge (P_{ki}), academic history (A_{Hi}), learning preferences (L_{Pri}), engagement patterns (E_{Pi}), assessment performance (A_{Pi}), and interaction logs (I_{Li}), into a unified learner profile that is continuously updated throughout the learning process. The learner profile is defined as,

$$LP_i = D_i, P_{ki}, A_{Hi}, L_{Pri}, E_{Pi}, A_{Pi}, I_{Li} \quad (1)$$

where, LP_i represents the learner profile. This representation consolidates multiple learner competencies into a unified learner profile, which provides a more holistic representation of learner competency than strategies that depend on limited academic indicators. The profile is continuously updated to reflect changes in the learner's progress, which is denoted as,

$$LP_i^{(t+1)} = LP_i^{(t)} \oplus L'_i \quad (2)$$

where, $LP_i^{(t)}$ represents the current learner profile, L'_i represents the newly produced learning evidence, and $\oplus$ represents the profile update operation. The newly obtained learning evidence consists of the assessment outcomes, engagement activities, and interaction records. Mathematically, the given formulation enables the learner profile to evolve by integrating newly generated learning evidence, thereby effectively capturing changes in learner competency, behaviour, and educational requirements throughout the learning process. Unlike static learner representations that mainly serve as descriptive inputs, the updated profile is then used as a decision-based state for KSE, which supports the estimated learner state to influence downstream personalized educational decisions. By combining heterogeneous learner information with a dynamic profile updating strategy, the presented subsystem overcomes the challenges of static learner modeling and restricted academic indicators.

C. Knowledge State Estimation Module

Following the student profile construction component, the KSE module is initiated to deliver an accurate and continuously updated representation of learner competency, which forms the basis for personalized learning decisions within the LITE-Trust framework. Previous AI-enabled learning systems mainly evaluate the learner's knowledge using static valuation scores or inadequate academic indicators, which makes them less efficient in extracting real-time learning development, detecting hidden learning gaps and accurately evaluating learner readiness for following concepts. To address these issues, the presented subsystem introduces a lightweight KSE model that employs the dynamically constructed learner profile together with multiple learner behaviour indicators, such as quiz performance, response time, learning activity patterns, revision frequency, and interaction history, to deliver a more reliable assessment of learner competency while preserving computational efficiency. The learner's knowledge state is denoted as a concept master vector, which is defined as,

$$KS_i = [ks_{i1}, ks_{i2}, \dots, ks_{in}], \quad 0 \leq ks_{ij} \leq 1 \quad (3)$$

where, KS_i represents the knowledge state of the learner i , and ks_{ij} denotes the mastery level of the concept j , and n represents the number of learning concepts. This representation models learner competency at the concept level, which allows accurate detection of mastered and weak concepts instead of depending on a single assessment score. The mastery level of each concept is defined as,

$$w_{QP} + w_{RF} + w_{RT} + w_{IH} = 1 \quad (4)$$

Subject to

$$w_{QP} + w_{RF} + w_{RT} + w_{IH} \geq 0 \quad (5)$$

where QP, RF, RT, and IH represent the quiz performance, revision frequency, response time, and interaction history, respectively, while w_1, w_2, w_3, w_4 represent their respective weights. To ensure a consistent combination of heterogeneous learner indicators, the quiz performance (QP), revision frequency (RF), response time (RT), and interaction history (IH) are first normalized to a common range of [0,1]. Quiz performance is directly represented as a normalized score, whereas revision frequency and interaction history are scaled using min-max normalization. A lower response time generally indicates more efficient learner interaction; response time is normalized using an inverse min-max transformation. The normalized indicators are therefore defined as,

$$QP_{\{i,c\}}^{\{*\}} = QP_{\{i,c\}} \quad (6)$$

$$RF_{\{i,c\}}^{\{*\}} = \frac{RF_{\{i,c\}} - RF_{min}}{RF_{max} - RF_{min}} \quad (7)$$

$$RT_{\{i,c\}}^{\{*\}} = 1 - \frac{RT_{\{i,c\}} - RT_{min}}{RT_{max} - RT_{min}} \quad (8)$$

$$IH_{\{i,c\}}^{\{*\}} = \frac{IH_{\{i,c\}} - IH_{min}}{IH_{max} - IH_{min}} \quad (9)$$

Here, $QP^{\{*\}}, RF^{\{*\}}, RT^{\{*\}}$ and $IH^{\{*\}}$ denote the normalized indicators. The inverse transformation for response time ensures that a shorter response time corresponds to a higher normalized behavioural score. To provide a reproducible weighting scheme, the KSE module assigns

$$w_{\{QP\}} = 0.40, w_{\{RF\}} = 0.20, w_{\{RT\}} = 0.15, w_{\{IH\}} = 0.25 \quad (10)$$

$$w_{\{QP\}} + w_{\{RF\}} + w_{\{RT\}} + w_{\{IH\}} = 1 \quad (11)$$

The concept-level mastery score is then calculated as,

$$M_{\{i,c\}} = w_{\{QP\}}QP_{\{i,c\}}^{\{*\}} + w_{\{RF\}}RF_{\{i,c\}}^{\{*\}} + w_{\{RT\}}RT_{\{i,c\}}^{\{*\}} + w_{\{IH\}}IH_{\{i,c\}}^{\{*\}} \quad (12)$$

Here, $M_{\{i,c\}}$ represents the estimated mastery of the learner for the concept. The larger weight assigned to quiz performance reflects its direct relationship with demonstrated academic competency, while revision frequency, response behaviour, and interaction history provide complementary behavioural evidence. The resulting normalized and weighted score is subsequently used to construct the learner knowledge-state vector and readiness estimate. The weighted formulation incorporates both behavioural and academic learning evidence, which provides a more accurate evaluation of concept mastery than conventional score-based approaches. By incorporating the updated learner model that contains new evidence, the evaluated knowledge state can reflect changes in learner competency over time. The learner readiness score is calculated as,

$$LR_i = \frac{1}{n} \sum_{j=1}^n ks_{ij} \quad (13)$$

where, LR_i represents the learner readiness score, ks_{ij} represents the mastery level of concept j , and n represents the number of concepts. The highest value of the readiness score

demonstrates stronger prerequisite knowledge and greater readiness for following concepts. Unlike learner models that demonstrate knowledge using a single aggregate performance value, the presented KSE generates a concept-level state that can differentiate individual mastery differences over concepts and thus provide more specific information for following educational decisions. Furthermore, a knowledge graph captures conceptual relationships and concept dependencies, which enables the system to detect knowledge gaps in prerequisite concepts and evaluate learner readiness using structured knowledge reasoning. By integrating behavioural learning analytics from the learner profile with a computationally efficient architecture for knowledge graph-based reasoning. The proposed subsystem provides a continuously updated and trustworthy representation of learner knowledge. The derived concept mastery vector, learner readiness score, and identified learning gaps are then provided to ELRA as learner-state variables for educational refinement. These outputs define the personalized search space and support the optimization of academic decision variables such as AI agent selection, prompt template, learning sequence, difficulty of assessment, feedback depth, and computation resource allocation. This enables more accurate personalization and educational decision-making while preserving computational efficiency.

D. Educational Library and Readers Optimization (ELRA)

Following the KSE subsystem, the assessed learner knowledge state derived from the continuously updated learner profile is transferred to the ELRA subsystem for personalized educational decision-making. The KSE is used to estimate the students' current knowledge level and learning needs. ELRA detects the most appropriate combination of AI agents, prompts, learning sequence, valuation difficulty, feedback quality, and computational resources. Thus, KSE estimates the learners' current requirements, whereas ELRA establishes the main optimization strategy that determines how those requirements should be addressed through an appropriate educational configuration. Table III explains the purpose and novelty of the ELRA module, and Fig. 2 represents the flowchart of the ELRA subsystem.

ELRA improves the LRA, which captures information-seeking behaviour through exploration, peer-based guidance, and exploitation. However, the traditional LRA is a versatile optimizer and does not obviously integrate learner knowledge state, educational trust, and explainability. ELRA addresses these issues by introducing learner-state-guided exploration, trust-aware selection, dynamic exploitation, learner-specific configuration, and multi-objective optimization while retaining the basic LRA search principle.

1) *Motivation*: The selection of ELRA is guided by the precise requirements of personalized academic optimization. Traditional metaheuristic algorithms, such as GA, PSO, and other population-based methods, are primarily designed for general-purpose numerical optimization and do not naturally represent learner profiles, educational trust, or learner-oriented objectives. Alternatively, LRA provides a lightweight exploration-exploitation strategy motivated by information-searching behaviour, making it an appropriate basis for

educational optimization. However, LRA does not obviously integrate learner state data or educational consistency. ELRA thus changes LRA to educational optimization by conditioning the search on the KSE-guided learner state, integrating trust into candidate evaluation and acceptance, adapting exploration to learner requirements, and collectively optimizing learning accuracy, explainability, trust, latency, and computational efficacy. This empowers ELRA to identify personalized and trustworthy educational configurations instead of generic optimization solutions.

2) *Educational representation and initialization*: The learner knowledge state gained from KSE is integrated into each academic configuration as,

$$I_i = [LP_i, KS, A_i, P_i, L_i, D_i, R_i, F_i] \quad (14)$$

where, LP_i represents the dynamic learner profile, KS represents the assessed learner knowledge state, while $A_i, P_i, L_i, D_i, R_i, F_i$ represents the AI agent, template of prompt, learning sequence, difficulty of assessment, computational resources, and feedback depth, respectively. These variables constitute the decision dimensions optimized by ELRA. In contrast to traditional optimizers, this representation explicitly relates the optimization variables to the learner's estimated state, which enables ELRA to search for personalized instead of standard educational configurations. This learner-state conditioning creates the search spaces as an education-specific decision space, in which optimal configurations are estimated based on the requirements of the current learner instead of general-purpose optimization solutions. The initial candidate population is defined as,

$$I_i^{(0)} = lb + r_i(ub - lb) \quad (15)$$

where, ub and lb represent upper and lower bounds and $r_i \sim u(0,1)$. This allocates candidate configurations over the search space and maintains diversity during initialization.

3) *KSE-guided knowledge exploration*: To adapt the exploration to educational requirements, ELRA changes the peer-guided LRA search using the learner state and trust data,

$$I_i^{P1} = I_i + \alpha_t(1 - 2r)(I_j - I_i) + \delta KS + \gamma(T_j - T_i) \quad (16)$$

where, I_i^{P1} represents the exploratory candidate configuration generated for the candidate i , I_i represents the current candidate configuration, I_j is a peer candidate reader, $r \sim u(0,1)$, T_j, T_i represents the trust scores of the current and peer configurations, α_t regulates the exploration magnitude, δ regulates the learner knowledge state influence, and γ denotes the trust-based exploration coefficient, which controls trust-based movement.

The conventional LRA update, which primarily uses peer-candidate information to determine the search direction, is augmented by the proposed ELRA update introduces two additional sources of optimization guidance: the learner knowledge state and the trust difference between candidate configurations. The learner-state term biases exploration toward configurations that address the learner's identified knowledge gaps and learning requirements, whereas the trust

term discourages movement toward configurations with lower educational reliability. The search direction is no longer determined solely by the geometric relationship between candidate solutions. Instead, ELRA produces a learner- and trust-conditioned search trajectory. This modification is algorithmic because it changes the candidate-generation rule itself and directly affects which configurations are explored during optimization.

Unlike the traditional LRA update, the presented formulation directs exploration based on both learners' needs and trust differences, instead of relying solely on peer information. This change prevents the exploration process from treating all candidate configurations equally and enables the search direction to remain stable with both the learners' estimated needs and the trustworthiness of the educational configuration.

4) *Trust-aware evaluation*: Every explored configuration is estimated using the trust value of the candidate configuration, which is defined as,

$$T_i = w_1 FC_i + w_2 + w_3 C_i + w_4 (1 - HP_i) + w_5 EQ_i + w_6 EC_i \quad (17)$$

where, $FC_i, CA_i, C_i, HP_i, EQ_i$, and EC_i represents factual correctness, curriculum alignment, confidence, hallucination probability, explanation quality, and educational completeness, respectively, with $\sum_{i=1}^6 w_i = 1$. A candidate undergoes refinement only when $T_i \geq T_{min}$. This trust limitation introduces a quality gate into the optimization process, which prevents candidates with unreliable or low-quality educational reliability from proceeding to refinement, instead of enabling fitness enhancement alone to determine candidate progression.

5) *Trust and state-guided adaptive exploitation*: For trusted configurations, ELRA performs adaptive optimization as

$$I_i^{P2} = I_i + \varepsilon_t (1 - 2r) \frac{\|ub-lb\|}{dt} + \omega(T_j - T_i) + \gamma(KS - EC_i) \quad (18)$$

where, I_i^{P2} represents the optimized candidate configuration, I_i represents the current candidate, ε_t regulates the exploitation step, d represents the decision variables count, t denotes the iteration number, r denotes a random variable; it is retained in the equation, ub, lb denotes the upper and lower bounds, γ denotes the learner-state influence coefficient, and EC_i represents the knowledge state needed to be associated with the current configuration. The gradually decreasing step size allows broader exploration in early iterations and fine-tuning in later iterations. The trust and state terms limit the optimization direction based on configuration reliability and learner-specific needs. This, while retaining the regular exploitation principle, is retained from LRA; ELRA conditions the optimization process on the evaluated knowledge state and configuration trustworthiness instead of depending only on numerical fitness improvement.

6) *Trust-aware acceptance and convergence*: A candidate substitutes the current configuration only when it meets trust, fitness, and learner-specific conditions,

$$I_i(t+1) = \begin{cases} I_i^{P2}, & T_i^{P2} \geq T_{min} \wedge F_i^{P2} < F_i \wedge \mathcal{A}(KS, I_i^{P2}) > \mathcal{A}(KS, I_i) \\ I_i, & \text{otherwise} \end{cases} \quad (19)$$

where, $I_i(t+1)$ represents the updated candidate configuration, I_i^{P2} denotes the optimized candidate T_i^{P2} represents the optimized candidate trust score, T_{min} denotes the minimum acceptable trust threshold, F_i^{P2} represents the optimized candidate fitness. ELRA minimizes the fitness function $F_i^{P2} < F_i$ indicates an objective improvement. $\mathcal{A}(KS, I_i)$ calculates the alignment between the learner's knowledge state and academic configuration. $\wedge$ represents logical AND, requiring all three conditions simultaneously. Therefore, ELRA accepts a promising solution only if it meets the trust requirements and improves the overall objective, and achieves better alignment with the learner state. This acceptance strategy prevents a candidate from being selected solely because of numerical objective enhancement, thus maintaining learner-specific suitability during the final solution-selection stage. The process ends when the highest iteration count is reached, or the enhancement becomes negligible. To ensure mathematical consistency and reproducibility, all learner-state, trust, alignment, and continuous configuration variables used by ELRA are represented in compatible normalized spaces. Let the learner knowledge state be represented by the concept mastery vector.

$$K_i = [k_{\{i1\}}, k_{\{i2\}}, \dots, k_{\{ic\}}]^{[T]} \in [0,1]^C \quad (20)$$

Here, C denotes the number of learning concepts and $kick_{\{ic\}}$ represents the estimated mastery of learner i for concept c . The corresponding learner-state vector used by ELRA is a C -dimensional normalized vector. An educational configuration is represented by

$$X_i = [a_i, p_i, l_i, d_i, r_i, f_i]^{[T]} \quad (21)$$

Here, a_i denotes the selected AI-agent configuration, p_i denotes the prompt-template configuration, l_i denotes the learning-sequence configuration, d_i denotes assessment difficulty, r_i denotes computational-resource allocation, and f_i denotes feedback depth. For numerical compatibility, continuous decision variables are normalized to $[0,1]$ as

$$x_{\{ij\}}^{\{*\}} = \frac{x_{ij} - x_j^{min}}{x_j^{max} - x_j^{min}} \quad (22)$$

Here, x_j^{max}, x_j^{min} denote the lower and upper bounds of the j -th decision variable. Categorical variables, such as agent and prompt selections, are represented by their predefined valid indices or one-hot configurations and are mapped back to their corresponding discrete values after optimization. The exploration and exploitation operations are performed in a common normalized decision space. The learner-state influence is incorporated through a dimension-compatible mapping

$$g(K_i, X_i) = W_{KK_i} - W_{XX_i} \quad (23)$$

Here, W_{KK_i}, W_{XX_i} are dimension-matching projection matrices that map the learner-state and configuration

representations into the same d-dimensional optimization space. The trust score is normalized to the interval [0,1] and is calculated as:

$$T(X_i) = w_F F_i + w_C C_i + w_Q Q_i + w_E E_i + w_H(1 - H_i) + w_G G_i, \sum_m w_m \geq 0 \quad (24)$$

Here, F_i, C_i, Q_i, E_i, H_i and G_i represent factual correctness, curriculum alignment, confidence, explanation quality, hallucination probability, and educational completeness, respectively. The transformation $(1 - H_i)$ converts hallucination probability into a positive reliability measure. The learner-configuration alignment is similarly normalized as

$$A(K_i, X_i) = 1 - \frac{\|g(K_i, X_i)\|_2}{\sqrt{d}} \quad (25)$$

A larger value of A indicates stronger compatibility between the learner's estimated knowledge requirements and the candidate educational configuration. Using these normalized quantities, the ELRA fitness function is expressed as

$$J(X_i) = w_L L(X_i) + w_T T(X_i) + w_A A(K_i, X_i) - w_R R(X_i) - w_C C(X_i) \quad (26)$$

Here, L denotes learning-performance utility, T denotes trust, A denotes learner-configuration alignment, R denotes response latency, and C denotes computational cost. All terms are normalized to [0,1], and the weights satisfy

$$w_L + w_T + w_A + w_R + w_C = 1 \quad (27)$$

The ELRA optimization with best configuration is consequently defined as

$$X^* = \operatorname{argmax}_{X \in \Omega} J(X) \quad (28)$$

This formulation establishes a common normalized representation for learner state, educational configuration, trust, and alignment and explicitly defines the feasible search space, objective function, acceptance conditions, and convergence criterion.

7) *Multi-objective fitness function*: ELRA collectively optimizes learning accuracy, trust, explainability, latency, and computational cost.

$$F(I) = \aleph_1(1 - L_A) + \aleph_2(1 - \mathcal{T}) + \aleph_3(1 - E) + \aleph_4 L + \aleph_5 C_C \quad (29)$$

where, $L_A, \mathcal{T}, E, L$, and C_C represents normalized learning accuracy, normalized trust, normalized explainability, normalized latency, and normalized computational cost, respectively, and $\aleph_1, \dots, \aleph_5$ denotes the objective weights, $\sum \aleph_i = 1$. In order to compare different objectives with each other, all five objectives were normalized to a scale of [0,1]. The benefit objectives, including learning performance, trust, and explainability, were normalized using min-max normalization as $X'_i = (X_i - X_i^{\min}) / (X_i^{\max} - X_i^{\min})$. The cost objectives, including latency and computation cost, were inversely normalized as $X'_i = (X_i^{\max} - X_i) / (X_i^{\max} - X_i^{\min})$.

In addition to modifying candidate exploration, ELRA changes the candidate evaluation and acceptance process. For a

candidate configuration XX , the optimization objective considers learning performance, trustworthiness, explainability, latency, and computational efficiency. This constrained acceptance mechanism distinguishes ELRA from conventional LRA, in which candidate evolution is primarily governed by the search and fitness mechanism. In ELRA, an improvement in the numerical objective alone is insufficient when the candidate does not satisfy the required trust and learner-alignment conditions. Thus, higher normalized values indicate better performance for all objectives, while the original cost measures are minimized. Afterwards, the normalized objectives were aggregated using the weighted objective configurations in Table XX, where the weighting groups represent different priorities among learning performance, trust, and efficiency. This design supports the sensitivity analysis to examine how changes in objective priorities affect the overall fitness and learning performance. The optimal configuration is obtained as,

$$I' = \operatorname{argmin}_I F(I) \quad (30)$$

Subject to

$$\mathcal{T}(I) \geq \mathcal{T}_{\min} \quad (31)$$

where, $\mathcal{T}_{\min} = 0.70$. This mathematical representation balances personalized learning efficiency, educational consistency, interpretability, response effectiveness, and resource consumption, thus treating educational optimization as a multi-criteria decision problem instead of relying solely on learning performance. In contrast to the conventional LRA objective, which is not formulated around these academic criteria, ELRA assesses promising configurations based on a unified educational objective that collectively includes effectiveness, interpretability, trustworthiness, computational efficiency and responsiveness.

The computational complexity is approximately $O[tn(F_C + T_C)]$, where t denotes the iteration count, n represents population size, F_C represents fitness cost, and T_C represents trust cost. The additional trust and learner-state evaluations introduce minimal computational cost while maintaining the lightweight population-based structure of LRA. Hence, ELRA is not only a relabeling of LRA; its exploration, optimization, and acceptance strategy explicitly integrate learner-state and trust information to generate personalized education configurations. The conventional LRA performs exploration and exploitation primarily through peer-guided candidate movement within a general optimization space. Its search behavior is determined by the current candidate, a peer candidate, and the corresponding search coefficients. In contrast, ELRA retains the basic exploration-exploitation principle of LRA but introduces learner-state conditioning, trust-aware candidate evaluation, and learner-specific acceptance constraints. Therefore, ELRA modifies the search dynamics rather than merely applying LRA to an educational task. ELRA extends this formulation by conditioning the candidate movement on the learner knowledge state and the trust difference between configurations. The search direction in ELRA is jointly determined by peer information, learner requirements, and trustworthiness. ELRA introduces a trust-

gated refinement mechanism. A candidate configuration is permitted to enter the subsequent optimization stage only when

$$T(X_{it}) \geq \tau_T, T(X_i^t) \geq \tau_T \quad (32)$$

Here, τ_T is the predefined trust threshold. The accepted candidate is then evaluated according to both optimization

quality and learner-configuration alignment. This introduces a constraint that is absent from the conventional LRA search process. The principal algorithmic distinction of ELRA is the transformation of an unconstrained peer-guided search into a learner-conditioned, trust-gated, and multi-objective educational optimization process.

TABLE III. PURPOSE AND NOVELTY OF ELRA MODULE

ELRA Module	Conventional LRA	Proposed ELRA Extension	Technical implication
Educational reader	General candidate/search representation	It represents AI agent, prompt, learning, feedback, assessment, and resources	Generates a learner-specific academic-search space
Knowledge exploration	Peer-guided exploration based on information-seeking patterns	Exploration is driven by the KSE-guided learner knowledge state and trust in information	Guides the search toward configurations aligned with the learners' current requirements
Trust evaluation	Candidate quality is mainly established through optimization fitness	Evaluates factual correctness, hallucination, curriculum alignment, explanation quality, and educational completeness	Presents educational trustworthiness in candidate evaluation
Candidate selection	Candidate evolution follows the traditional search and fitness strategies	Candidates should satisfy a minimum trust threshold before further optimization	Prevents untrustworthy educational configurations from entering additional optimization
Adaptive exploitation	Traditional exploitation of promising search regions	Exploitation is dynamically guided by learner state, trust, and iteration-dependent step size	Guides fine-tuning toward learner-aligned and reliable configuration while gradually reducing the search step
Candidate acceptance	Selection is mainly based on the enhancement within the optimization process	Acceptance includes trust, objective enhancement, and learner-configuration alignment	Maintains both solution quality and learner-specific fitness
Objective formulation	General optimization objective	Collectively optimizes learning accuracy, explainability, trust, latency, and computational cost	Adapts the search into an academic-specific multi-objective optimization problem
Final solution	Generally optimized candidate	Trustworthy configuration aligned with the learner's knowledge state and academic requirements	Generates a learner-specific, resource-aware, and trustworthy academic configuration for personalized decision-making

Pseudocode of ELRA module

Input: learner profile L_p, lb, ub , knowledge state KS , population size n , maximum iteration T_{max} , trust threshold T_{min} , and convergence tolerance ε

Output: optimal educational configuration I'

Initialize population $I_i^{(0)}$ using Eq. (8) $i = 1, \dots, n$

Evaluate fitness $F(I_i)$ using Eq. (13) and trust score $\mathcal{T}_i$ using Eq. (10) for each I_i

Select the optimal trusted configuration I'

For $t=1$ to T_{max} do

For $i=1$ to n do

Select optimal configuration I_j , where $j \neq i$

Generate experimental configuration I_i^{P1} using Eq. (9)

Compute trust score $\mathcal{T}_i^{P1}$ using Eq. (10)

If $\mathcal{T}(I) \geq T_{min}$ then

Generate optimized configuration I_i^{P2} using Eq. (11)

Compute fitness $F(I)$ using Eq. (13), trust score and $\mathcal{T}_i$ using Eq. (10), and learner-configuration alignment

If $\mathcal{T}(I) \geq T_{min}$ the learner-state requirements satisfy Eq. (11), then

Update as the trusted candidate with minimum fitness

End if

End if

End for

Update the population toward the trusted candidate using Eq. (9)

If the convergence condition is satisfied, then break

End for

Return I'

E. Lightweight Multi-Agent Coordinator

Following the KSE and ELRA subsystems, the optimized academic configuration produced by ELRA, such as selected AI agent prompt templates, learning sequence, assessment difficulty, feedback depth, and computational resource allocation, is directly fed into the activation strategy of the lightweight multi-agent coordinator. The multi-agent coordinator is responsible for deciding which agents are responsible for delivering it, which includes the tutor, assessment, feedback, recommendation, and academic integrity agents. However, all agents being active for each interaction would involve high computational costs and delays; the presented module introduces context-adaptive agent activation. The tutor agent offers explanations and guidance, the assessment agent computes learner performance, the feedback agent provides feedback to the learner, the recommendation agent selects relevant learning resources, and the academic integrity agent validates content and activities for academic integrity issues. Their specialized roles enable educational tasks to be dynamically allocated based on the learners' current requirements. Every educational agent is implemented as a task-specific role, based on a shared Generative Pre-trained Transformer 4 omni (GPT-4o) mini language model, rather than a separate model. Every agent takes the learner knowledge state, ELRA-optimized configuration, and task-specific context as inputs, and outputs role-specific responses, such as explanations, evaluations, feedback, learning-resource recommendations, or integrity checks, respectively. The coordinator calculates the relevance of every agent from the

current learning context and learner-state/ELRA configuration, activates the agents whose relevance scores exceed the activation threshold, and respects the quality and computational-cost constraints (Eq. (16)-(18)). The relevance score is determined according to the compatibility between the current academic task, the KSE-derived learner knowledge state, and the ELRA-optimized academic configuration. Thus, agent activation is learner- and task-dependent instead of based on a derived execution sequence. The active agent set is defined as,

$$AG_t = \{a_j | r_j(KS, I', C_{Et}) \geq \theta_j\} \quad (33)$$

where, AG_t represents the set of active agents at iteration t , a_j represents the j -th educational agent, $r_j(\cdot)$ represents the relevance score of agent j , C_{Et} represents the current learning context, I' denotes the ELRA-optimized educational configuration, and θ_j denotes the activation threshold. Thus, agents are activated only when their relevance scores exceed the specified threshold. This KSE-and ELRA-based selection minimizes unnecessary agent execution and identifies the gap in redundant multi-agent processing and maximizes response latency. The selected agent outputs are denoted as,

$$OP_t = Cf(\{OP_j : j \in AG_t\}, KS, I') \quad (34)$$

where, OP_t represents an active agent j output and $Cf(\cdot)$ represents the agent coordination function. By integrating KS and I' , the coordinator produces agent outputs based on the learners' state and personalized educational requirements, thus identifying isolated agent operation and unreliable responses. To reduce agent utilization while preserving educational quality, the optimal active set is defined as,

$$AG'_t = \arg \min_{AG_t} [\lambda_c |AG_t| + \lambda_l L_t] \text{ s.t. } RQ(AG_t, KS, I') \geq RQ_{min} \quad (35)$$

where, $|AG_t|$ represents the active agent count, λ_c , and λ_l represents the weight values of computational cost and response latency. The coefficients satisfy $\lambda_c, \lambda_l \geq 0, \lambda_c + \lambda_l = 1$. L_t denotes response latency, and RQ denotes educational response quality, and RQ_{min} denotes the minimum acceptable response quality.

This mathematical representation identifies the minimum set of agents that meets the required quality, which provides the major contribution of lightweight, quality-based multi-agent implementation. Thus, agent selection is formulated as a

forced active -set decision that collectively considers educational quality and processing cost, instead of evaluating the relevance of each agent individually. Overall, the subsystem contributes learner-based, ELRA-driven, dynamically activated, and quality-based multi-agent coordination, which addresses redundant computation, latency, isolated agent operation, and unreliable educational responses. Table IV emphasizes the practical contribution of every specialized agent and its role in the designed framework. In contrast to traditional multi-agent systems that mainly combine individual agents, the presented coordinator assigns learner-personalized responsibilities and adaptively activates agents based on the KSE-based state and ELRA-optimized configurations.




Fig. 2. Flowchart of ELRA module.

TABLE IV. CONTRIBUTION OF MULTI AGENTS

Agents	Proposed Contribution
Tutor Agent	Provides explanations and guided learning based on the KSE-based learner state, which enables personalized knowledge support.
Assessment Agent	Produces and evaluates assessments using the ELRA-adapted difficulty, which allows learner -adaptive assessment.
Feedback Agent	Delivers targeted corrective and formative feedback with an appropriate level of depth, which supports knowledge-gap-specific learning support.
Recommendation Agents	Chooses learning resources and activities based on KS , and I' , which enables optimization-based resource personalization.
Academic Integrity Agent	Evaluates academic content and activities before delivery, which provides trust-aware integrity verification throughout the learning pipeline.

F. Explainability Engine

Following the Lightweight Multi-agent coordinator, the outputs of the selected agents OP_t along with the KSE-based learner state KS_i and ELRA-optimized configuration I' are offered to the explainability engine. While personalized AI decisions can enhance learning outputs, their limited transparency may reduce trust among learners and instructors. Thus, the presented engine produces explanations, reasoning, and confidence for each academic decision, which identifies the interpretability gap in AI-based personalized education. The explanation is expressed as,

$$EX_t = \varepsilon(KS_i, I', OP_t, E_{ct}) \quad (36)$$

where, EX_t represents the produced explanation, $\varepsilon(.)$ represents the explainability function, I' represents the ELRA-optimized configurations, and E_{ct} represents the learner's current educational context. The formulation is conditioned on the learner state, optimized configurations, and agent outputs, explaining why a concept was learned, why assessment difficulty was chosen, why a revision action was taken, or why a learning resource was recommended, thus addressing the issue of context-free recommendations. The decision reasoning and confidence are obtained as,

$$DR_t = DR(KS_i, I', OP_t), S_t = S(DR_t) \quad (37)$$

where, DR_t represents decision reasoning and S_t represents the respective confidence score. $DR(.)$ represents the reasoning generation function, $S(.)$ represents the confidence estimation function; both quantities are derived from the learner state, optimized configuration, and coordinated agent output. Hence, the engine explains the rationale for each decision and its confidence level, thereby enhancing educational transparency and instructor confidence. The final explainable output is formulated as,

$$Z_t = \{EX_t, DR_t, S_t\} \quad (38)$$

where, Z_t collectively represents explanation, reasoning, and confidence. This formulation connects the generated explanation to the information that contributed to the educational decision, which facilitates the rationale to reflect the learner-specific state and optimal educational configuration instead of generating individual generic explanations. Hence, the presented engine provides learner-state-aware and ELRA-based explanations, reasoning, and confidence, thereby bridging the transparency and trust gap in personalized academic decision-making.

G. Trust Evaluation Engine

Following the explainability engine, the trust evaluation engine executes final trust validation of the educational response produced by the optimal agents. While ELRA already computes trust during optimization, that estimation is used to select optimal candidate configurations. Unlike the presented trust evaluation engine, which validates the actual produced response prior to delivery, it addresses the gap between trustworthy configuration selection and trustworthy final academic output. The final response trust score is expressed as,

$$T_{St} = w_{fc}FC_t + w_{ca}CA_t + w_h(1 - H_t) + w_cC_t \quad (39)$$

Subject to:

$$w_{fc} + w_{ca} + w_h + w_c = 1 \quad (40)$$

where, T_{St} represents the final response trust score of the generated response, FC_t represents factual correctness, CA_t represents curriculum alignment, H_t represents hallucination probability, and C_t denotes confidence. While w_{fc}, w_{ca}, w_h, w_c represents their respective weights. The term $1 - H_t$ transforms the estimated hallucination probability into a respective reliability measure. Hence, the mathematical representation collectively evaluates factual reliability, curriculum suitability, hallucination risk, and confidence, instead of depending solely on model-generated confidence. A response is accepted only when

$$T_{St} \geq T_{Smin} \quad (41)$$

where, T_{St} represents the final response trust score, and T_{Smin} represents the minimum acceptable response threshold. All trust components are normalized to $[0,1]$, making $T_{St} \in [0,1]$. A response satisfying $T_{St} \geq T_{Smin}$ is accepted for delivery, whereas a response below this threshold is rejected or refined before delivery. In the experimental configuration, the trust threshold is fixed at 0.70. For this final step, the produced response is evaluated against both the corresponding educational knowledge and reference information for the associated learning task. Factual correctness is established by comparing the derived responses to the provided reference information. Meanwhile, the degree of alignment with the curriculum is determined by the target concept and its requirements. Thus, hallucination probability is defined as the ratio of inaccurate information that was found in comparison with the reference information during fact-checking. The confidence score is extracted from the model confidence assessment. Finally, these metrics are used to inform the response-level acceptance decision. The threshold therefore functions as a response-level acceptance limit, which makes trust evaluation an active control strategy instead of only a post-hoc reliability indicator. Hence, the presented engine acts as a final trust-aware integrity control, which prevents undependable educational responses from reaching the learner. Unlike ELRA, in which trust drives the candidate configuration selection and optimization, this component analyses the final agent's response. This two-stage trust strategy boosts the developed framework by incorporating both learning processes: the trust-aware educational optimization and final fact-checking step, which improves the accuracy of hallucination-free and curriculum-aligned educational content while enabling more reliable personalized decision-making.

H. Academic Integrity Analyzer

Following the trust evaluation engine, the academic integrity analyzer estimates whether AI-assisted learning maintains originality, critical thinking, independent problem-solving, and authentic learning behaviour. Traditional integrity methods focus primarily on plagiarism, which may overlook excessive AI dependency and reduced learner independence. Hence, the introduced analyzer aims to mitigate these issues by monitoring the authenticity and adequacy of AI-assisted learning. The educational integrity score (I_{At}) is expressed as,

$$I_{At} = w_1 OR_t + w_2 CT_t + w_3 PS_t + w_4 AL_t \quad (42)$$

Subject to

$$w_1 + w_2 + w_3 + w_4 = 1 \quad (43)$$

where, OR_t represents originality, CT_t represents thinking, PS_t represents independent problem-solving, and AL_t represents authentic learning behaviour, each estimated from relevant learner interaction and performance patterns available in EdNet and OULAD. $w_1, \dots, w_4$ represents corresponding weight values $w_k \geq 0 \sum w_k = 1$. The weighted formulation integrates multiple behavioural indicators, thereby delivering a more holistic integrity assessment than plagiarism detection alone. AI dependency is evaluated as,

$$AD_t = \frac{\zeta_{AI}}{\zeta_{AI} + \zeta_I} \quad (44)$$

where, AD_t represents the estimated AI-dependency score, ζ_{AI} represents AI-assisted interactions, and ζ_I represents independent learning interactions. A maximum AD_t represents greater AI dependency. Higher AD_t indicates greater reliance on AI-assisted interactions. The integrity need is satisfied when $AI_t \geq AI_{min}$, $AD_t \leq AD_{max}$ where AI_{min} represents the minimum integrity score and AD_{max} represents the maximum acceptable AI dependency. These constraints ensure that AI assistance does not take over from genuine human learning. The indicators identified by the academic integrity analyzer should be regarded as behavioural estimates rather than measures of academic integrity. In fact, EdNet and OULAD provide information on the learning process and its outcomes, from which integrity and the extent to which students rely on AI are inferred. Put another way, since the input data cannot provide ground truth about academic integrity and AI reliance, the analyzer does not claim to directly assess these aspects, but rather it identifies learning behaviours that may indicate academic integrity issues and/or reliance on AI. Therefore, the contribution of the proposed work is an AI-aware, behaviour-based integrity check that builds upon the preceding trust verification and aims to bridge the plagiarism-centric gap in academic integrity assessment.

The academic-integrity component is currently evaluated using behavioral indicators derived from learner interaction patterns rather than independently annotated ground-truth integrity labels. Therefore, its outputs should be interpreted as risk indicators for potential academic-integrity concerns, rather than definitive evidence of misconduct. The framework does not claim to establish academic-integrity violations solely from these behavioral estimates. Validation against expert annotations or independently labeled behavioral indicators is identified as an important direction for future evaluation.

I. Personalized Learning Dashboard

Following the academic integrity analyzer, the personalized learning dashboard introduces the learners' learning analytics, personalized learning recommendations, concept-level progress, knowledge gap alerts, trust alerts, and academic-integrity alerts using a lightweight interface. Traditional educational dashboards provide learning information to a learner in a non-adaptive manner. Instead, the presented dashboard employs KSE-derived learner state, ELRA-optimal

configuration, multi-agent outputs, trust assessment, and academic integrity results to provide an integrated learner view. The dashboard limits the information that is relevant to the learning context, such as the level of mastery, recommendations of resources, learning needs, assessment results, trust alerts, and academic-integrity alerts, thus enabling faster processing of the information and interaction.

The novelty of LITE-Trust lies in the explicit dependencies among the proposed methods rather than the individuality of each technique used. In particular, while the continuously updated learner profile informs KSE, KSE informs the learner's knowledge state for ELRA, the latter informing the educational configuration that determines the agent's selection. The response is then evaluated by the trust engine. As such, it can be argued that the mathematical formulations describe a learner-state-driven decision-making process where each step informs the next one rather than several independent procedures combined together.

The LITE-Trust framework is designed to perform lightweight execution through compact KSE, limited iterations of ELRA, selective agent activation, and optimized prompt and retrieval processes. First, KSE reduces subsequent operations by using a compact knowledge-state representation. ELRA further limits search by considering a fixed population size and iterations. Finally, dynamic coordination allows for activating only specific agents according to the incoming request. In addition, prompt optimization and dense-vector-retrieval reduce additional computational burdens. Hence, the efficiency gains mainly result from reduced representation, limited optimization step, selective execution, and retrieval-generation costs. The reported parameter count excludes the parameters of the external LLM employed by LITE-Trust.

J. Computational Complexity Analysis

The computational efficacy of LITE-Trust is achieved through the implementation of lightweight KSE representation, reduced ELRA search, optimal agent activation, and optimized prompt and retrieval processing. The overall computational complexity is then decomposed based on the major processing components of LITE-Trust. The KSE component requires $O(nl)$, where n represents the count of learner records, and l represents the average learner-interaction sequence length. This follows because each of the n learner records processes an average of l interaction elements. ELRA requires $O(tnF_c)$, where t denotes the iteration count, n represents population size, F_c represents fitness estimation cost for one candidate solution. Therefore, n candidates are evaluated during each of t iterations. Dynamic agent selection requires $O(V + k)$ in the worst case and $O(k)$ after selective activation, where V is the total number of available agents and k is the number of activated agents. Since LITE-Trust achieves only the agents required for a particular learner request, $k \ll V$ in a typical execution, thereby avoiding unnecessary agent computation and reducing inference latency and utilization. Trust evaluation requires $O(kq)$, where q represents the cost of computing the trust-related indicators. Explainability processing requires $O(e)$, representing the cost of generating and processing the corresponding explanations, where e denotes the number of explanatory factors. Finally, retrieval and generation require

$(\log N + r + C_g)$, where N represents the indexed knowledge-base size, r represents the number of retrieval items, and C_g represents generation cost. Consequently, the overall computational complexity is derived as the sum of the component complexities $C_{\text{LITE-Trust}} = O(nl + tnF_c + V + k + kq + e + \log N + r + C_g)$. This component-wise formulation reflects the actual operations performed by each module. The presented lightweight design reduces computational complexity through compact KSE representation, bounded ELRA search, selective agent activation, and optimized retrieval and generation, thereby reducing unnecessary computation and inference latency while preserving the required educational quality.

To provide a more realistic efficiency assessment, the computational analysis considers the complete execution pipeline rather than only the local algorithmic complexity. The evaluation accounts for GPT-4o mini API invocation, embedding generation, retrieval operations, multiple-agent execution, trust verification, and ELRA optimization iterations. For each request, end-to-end latency is considered together with local processing time, API response time, network communication overhead, and the number of activated agents. External inference and API usage costs are also recognized separately from the computational cost of the locally deployed components. This end-to-end assessment provides a more representative measure of the practical efficiency of LITE-Trust.

IV. RESULTS AND DISCUSSION

A. Experimental Environment and Datasets

The LITE-Trust framework is experimentally evaluated using two benchmark educational learning analytics datasets, such as OULAD and EdNet. These datasets provide the necessary learner interaction, assessment, and activity data required to evaluate the proposed learner knowledge estimation, personalized optimization, adaptive agent coordination, trust evaluation, and academic-integrity analysis modules. OULAD contains learner demographic information, valuation results, module information, and Virtual Learning Environment (VLE) interactions, which enable learner progress analysis and behavioural pattern analysis. EdNet provides large-scale learner-question interaction data, which includes question attempts and response histories, that supports concept-level knowledge estimation and personalized learning analysis.

The EdNet and OULAD primarily provide learner interaction, assessment, and educational-process information. These datasets are therefore used to construct learner profiles, estimate knowledge states, and evaluate personalized learning performance. For generated educational content, hallucination, trust, explainability, and academic-integrity evaluation, the dataset records are supplemented with corresponding reference answers, learning objectives, and required educational information where applicable. These references provide the basis for checking factual consistency, content relevance, explanation quality, and unsupported information. The evaluation criteria are applied consistently across all responses, while dataset-derived learner information is kept separate from the ground-truth information used for content-level validation.

TABLE V. DESCRIPTION OF THE OULAD AND EdNET DATASETS USED FOR TRAINING AND TESTING THE LITE-TRUST

Dataset	Category	Total	Training (80%)	Testing (20%)
OULAD	Students	32,593	26,074	6,519
	VLE interactions	10,655,280	8,524,224	2,131,056
EdNet	Students	784,309	627,447	156,861
	Interactions	131,441,538	105,153,230	26,288,307

Table V represents the dataset descriptions, which include total samples and training and testing samples. For both datasets, the data were split into 80% training and 20% test samples at the level of learners so that no interactions from the same learner would appear in both train and test sets. For the purposes of the experiments described in this study, the evaluation was carried out in five independent runs with random seeds 11, 22, 33, 44, 55, while preserving the same learner-level split protocol across all runs. The values reported for the experiments are the mean $\pm$ standard deviation over the five runs. The test set was kept hidden during training, parameter tuning, and optimization.

TABLE VI. HARDWARE AND SOFTWARE CONFIGURATION USED FOR EXPERIMENTAL EVALUATION

Hardware Configuration	
Processor	Intel Core i7-12700H, 14 cores
Storage	512 GB SSD
RAM	16GB
Software Configuration	
Operating system	Ubuntu 22.04 LTS
Programming Language	Python 3.9
Development environment	Pycharm
Embedding API	Text-embedding-3-small

TABLE VII. PARAMETER CONFIGURATION AND EXPERIMENTAL SETTINGS OF THE LITE-TRUST

Parameters	Value
LLM	GPT-4o-mini (external API)
Embedding model	Text-embedding-3-small (external API)
Agent activation	Context-adaptive
Tokenization	Standard tokenizer
Retrieval	Dense-vector similarity search
Trainable parameters	1.92 M (LITE-Trust components only)
Prompt strategy	Compact context-aware prompt
KSE weight constraint	$\sum_j w_j = 1$
ELRA population size	30
ELRA exploration coefficient	0.5
ELRA exploitation coefficient	0.5
ELRA maximum iterations	100
Trust threshold	0.70
Agent activation threshold	0.70
Convergence tolerance	0.001

Table VI represents the system configuration, and Table VII illustrates the parameter configuration of the LITE-Trust framework. The reported 1.92 M trainable parameters are for the locally deployed LITE-Trust components only; the externally accessed models (GPT-4o mini and text-embedding-3-small) are not included. The costs associated with the external LLM and embedding services, such as network/API latency and generated cost, are treated separately from the computational costs of the locally deployed LITE-Trust components. Here, the activation threshold was fixed at 0.70 across all experiments to maintain a stable relevance criterion for selective agent execution.

B. Evaluation Metrics

1) Educational performance

a) *Knowledge gain*: It is generally denoted by the improvement in the learner's knowledge from the pre-learning assessment to the post-learning assessment, which is expressed as,

$$G_K = \mathfrak{S}_{post} - \mathfrak{S}_{pre} \quad (45)$$

where, $\mathfrak{S}_{post}$ represents the pre-learning assessment score, and $\mathfrak{S}_{pre}$ represents the post-learning assessment score.

b) *Prediction accuracy*: It is given by the ratio of the number of true responses identified by the learners, which is defined as,

$$ACC = \frac{T_P + T_N}{T_P + T_N + F_P + F_N} \quad (46)$$

where, T_P, T_N, F_P, F_N represents true positive, true negative, false positive, and false negative, respectively.

c) *F1-Score*: It is used to measure the balance between precision and recall, which is expressed as,

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (47)$$

d) *Recommendation precision and recall*: Precision is used to measure the percentage of relevant recommendations among the selected resources, while recall is used to measure the percentage of relevant resources that were effectively recommended.

$$RP = \frac{T_{PR}}{T_{PR} + F_{PR}}, RR = \frac{T_{PR}}{T_{PR} + F_{NR}} \quad (48)$$

where, T_{PR}, F_{PR} , and F_{NR} represent recommended relevant resources, recommended irrelevant resources, and relevant but unrecommended resources, respectively.

e) *Student satisfaction*: It is used to measure the average learner satisfaction according to the adopted questionnaire or Likert-scale responses, which is defined as $S_S = \frac{1}{m} \sum_{i=1}^m \mathfrak{s}_i$. Where $\mathfrak{s}_i$ represents the learner satisfaction score, and m represents the participating learners' count.

2) Trustworthiness

a) *Trust score*: It is used to measure the overall trustworthiness of the produced educational outcomes by incorporating factual correctness, confidence, curriculum alignment, explanation quality, and educational completeness. The trust score is expressed as, $T_S = \sum_{l=1}^m w_l T_l$, $\sum_{l=1}^m w_l =$

1 Where T_l represents each trust-relevant measure and w_l represents its respective weight value.

b) *Hallucination Rate*: It is used to measure the percentage of produced responses containing unsupported or inaccurate information, which is expressed as,

$$H_R = \frac{n_h}{n_g} \times 100 \quad (49)$$

where, n_h represents the count of responses including hallucinated information and n_g represents the total count of generated responses. A reduced hallucination rate indicates the highest trustworthiness. To calculate the relative enhancement in hallucination control, the hallucination reduction is formulated as,

$$HR_{rd} = \frac{H_{R,base} - H_{R,LITE}}{H_{R,base}} \times 100 \quad (50)$$

where, HR_{rd} denotes the hallucination reduction, $H_{R,base}$ denotes the hallucination rate of the LLM-based tutoring baseline, and $H_{R,LITE}$ denotes the hallucination rate of LITE-Trust under the same evaluation protocol. Hence, the hallucination reduction demonstrates the relative decrease in hallucinated responses obtained by Trust compared with the selected baseline.

c) *Factual consistency*: Factual consistency refers to the percentage of responses that are factually consistent with educational knowledge or reference information, which is expressed as,

$$F_C = \frac{n_c}{n_v} \times 100 \quad (51)$$

where, n_c represents the count of statements that are factually consistent and n_v represents the total count of the assessed statements. The higher the factual consistency, the better the educational performance is achieved.

For the trust assessment, the EdNet and OULAD datasets provide the context of learner interaction, assessment, and the educational process, but are not used as direct grounds for answers. The produced responses are scored against the ground-truth answers and corresponding educational question-answer information provided with the evaluated learning tasks, which serve as reference information for factuality and hallucination evaluation. The factual correctness is determined by factual statements extracted from the generated responses and compared to the reference information. In addition, the alignment to the curriculum is evaluated based on the target concept and learning objective of the given prompt. The hallucination probability is inferred from the response's unsupported or incorrect statements identified against the reference information. The confidence score is obtained from the model confidence assessment, while the explanation quality is based on the response's relevance, completeness, clarity, and overall consistency. Finally, the educational completeness is defined by the response's ability to fully incorporate the required elements of the given prompt.

3) Explainability

a) *Explanation fidelity*: It is generally used to measure the accuracy with which the obtained explanation reflects the

factors leading to an educational decision. It is expressed as follows,

$$E_F = \frac{n_c}{n_e} \times 100 \quad (52)$$

where, n_c represents the number of explanations that are consistent with the decision-making process and n_e represents the total number of explanations evaluated.

b) Explanation quality: It is generally used to estimate the explanation quality according to its relevance, completeness, clarity, and consistency. $E_Q = \frac{1}{n} \sum_{i=1}^n q_i$, where q_i represents the quality score of the explanation i , and n denotes the count of evaluated explanations.

c) User interpretability: it is used to assess learners' ability to understand and interpret the obtained explanations, which is expressed as $I_U = \frac{1}{n} \sum_{i=1}^n u_i$, where u_i represents the interpretability rating generated by the learner i , typically assessed using a Likert-scale evaluation, where higher scores demonstrate greater user interpretability.

For the purpose of explainability evaluation, the fidelity of the produced explanation to the factors that contributed to the corresponding educational decision is determined. Explanation quality is measured in terms of relevance, completeness, clarity, and consistency of the derived explanation. User interpretability is then derived from the learner-oriented interpretability metric as specified in the evaluation procedure. The evaluation is conducted using the same criteria and scoring process for all obtained explanations; hence, no separate expert-panel or multi-rater evaluation is necessary, and inter-rater reliability is not determined. The fidelity, explanation quality, and interpretability measures reported in the study are automated, criterion-based indicators designed to provide a consistent and reproducible assessment of generated explanations. They quantify properties such as consistency with the generated response, relevance, completeness, clarity, and alignment with the associated learning content. These measures are therefore intended as proxy indicators of explanation quality.

4) Efficiency

a) Inference latency: It is used to measure the average execution time taken by the model to generate an output, which is formulated as,

$$L_I = \frac{1}{n} \sum_{i=1}^n t_i \quad (53)$$

where, t_i represents the inference time for interaction i .

b) Memory consumption: It is used to measure the memory consumed by the model during execution. Therefore, low memory utilization indicates more effective utilization of system resources and supports placement in resource-constrained educational environments.

c) Number of model parameters: It is used to estimate the total number of trainable parameters of the LITE-Trust model, which reflects the memory consumption and model complexity, where a lower parameter count indicates that the model is lightweight.

d) FLOPs: It is used to evaluate the number of floating-point operations required during model inference, which reflects the computational complexity of the LITE-Trust architecture, while fewer FLOPs indicate lower computational complexity.

e) Computational cost: It is used to measure the computational burden of the system, with lower values indicating more computational efficiency and lighter algorithms.

f) Response time: It is used to compute the overall processing time taken from receiving the request from the user until the moment of providing the final response, which is defined as,

$$R_T = t_o - t_i \quad (54)$$

where, t_o , and t_i represent the request delivery, and request response times, respectively.

To improve the objectivity and reproducibility of the trust model, all trust components are determined using predefined evaluation criteria and are kept fixed across the experiments. Factual correctness and hallucination are evaluated against the corresponding ground-truth answers and educational reference information, while curriculum alignment is assessed with respect to the target concept and learning objective of each task. Confidence is obtained from the model's confidence assessment. Explanation quality is evaluated according to relevance, completeness, clarity, and consistency, whereas educational completeness reflects whether the required instructional elements of the task are adequately covered. These criteria are applied uniformly to all evaluated responses rather than being adjusted according to the obtained performance. The weights assigned to the individual trust components are fixed before the experimental evaluation and are applied identically to all datasets, learner profiles, and generated responses. The weight selection is based on the relative importance of factual reliability, instructional relevance, and response quality for educational recommendation. No weight is modified based on the performance of an individual model or test sample. This fixed-weight protocol prevents post-hoc adjustment of the trust score and improves reproducibility across experiments.

The reliability of the trust assessment is further examined by evaluating the individual trust components separately from the overall trust score. Factual consistency and hallucination rate are evaluated against reference educational information, while curriculum alignment and educational completeness are assessed against the predefined learning objectives and task requirements. Explanation quality is evaluated independently using its relevance, completeness, clarity, and consistency criteria. The component-level results are then compared with the aggregated trust score to verify that high trust values correspond to responses with stronger factual reliability, lower unsupported content, appropriate curriculum coverage, and better educational completeness.

To improve the reliability of the knowledge-gain evaluation, the pre-learning and post-learning assessments are compared under consistent learner and assessment conditions.

The same learner profiles, target concepts, assessment criteria, and evaluation protocol are maintained across the compared methods. Learner characteristics and prior knowledge are considered when interpreting the observed improvement, while temporal progression, repeated exposure, and differences in learner participation are controlled as far as permitted by the available data.

C. Overall Performance Comparison

The LITE-Trust is compared with existing approaches such as conventional ITS [15], conventional multi-agent systems [10], LLM-based tutoring [6], non-optimized multi-agent systems [12], optimization-based approach [20], and adaptive learning models [7], to evaluate its performance. These baselines were selected as they represent major paradigms of the educational system addressed by LITE-Trust, namely, rule-based tutoring, multi-agent coordination, LLM-assisted learning, adaptive learning modelling, and optimization-driven personalization. Each baseline is evaluated on the same EdNet and OULAD datasets, with the same train/test partitions, hardware configuration, and evaluation metrics as LITE-Trust, using comparable learner-level partitions of the data to ensure validity of comparison, with variations specific to each method and their implementation details retained.

For a fair comparison, all baseline methods and the proposed LITE-Trust framework were evaluated under the

same experimental setting, including the same datasets, learner profiles, input information, evaluation protocol, and computing environment. The same experimental scenarios and stopping criteria were maintained for all methods. Each baseline was implemented using its defined methodological characteristics, while the proposed framework used the same available input and evaluation conditions. Model-specific parameters, prompts, and configuration settings were kept fixed within each method and were not selectively adjusted to favor the proposed approach. This ensures that the reported differences are primarily attributable to the methodological differences between the compared approaches rather than changes in the experimental conditions.

The results in Tables VIII–XV are presented as descriptive performance comparisons, with mean values and standard deviations used to indicate central performance and variability across experimental runs. Formal statistical testing was applied to the optimizer-level comparison because ELRA was specifically evaluated against alternative optimization strategies. Nevertheless, to strengthen the statistical reliability of the broader baseline comparisons, confidence intervals and/or appropriate pairwise statistical tests can additionally be reported for the headline metrics. Differences such as 95.6% versus 90.1% are interpreted as observed performance improvements rather than statistically significant differences unless supported by the corresponding statistical test.

TABLE VIII. ANALYSIS OF EDUCATIONAL PERFORMANCE ON EdNet AND OULAD DATASET

Model	EdNet			OULAD		
	KG (%)	Accuracy (%)	F1 (%)	KG (%)	Accuracy (%)	F1 (%)
Conventional ITS	12.7±0.5	85.5±0.5	85.8±0.5	11.9±0.6	85.1±0.5	85.1±0.5
Conventional multi-agent system	13.6±0.5	87.5±0.6	86.0±0.7	12.9±0.6	87.0±0.5	86.1±0.7
LLM-based tutoring	14.5±0.6	87.6±0.5	87.9±0.6	13.4±0.7	86.9±0.4	87.1±0.6
Non-optimized multi-agent system	15.2±0.4	89.1±0.5	88.7±0.5	14.2±0.6	88.8±0.6	88.3±0.6
Adaptive Learning model	15.3±0.7	89.2±0.5	89.2±0.5	14.5±0.5	88.9±0.6	88.3±0.5
Optimization-based approach	16.5±0.4	90.1±0.5	89.9±0.6	15.1±0.5	89.8±0.5	89.0±0.5
LITE-Trust	18.8±0.5	95.6±0.2	94.5±0.4	17.9±0.4	94.8±0.3	93.5±0.4

Table VIII explains the numerical results of educational performance on the EdNet and OULAD datasets. The LITE-Trust achieves the highest performance with an accuracy of 95.6%, an F1-score of 94.5%, knowledge gain of 18.8% on the EdNet dataset, and 94.8%, 93.5%, and 17.9% on OULAD, respectively. The optimization-based approaches achieve only 90.1% accuracy and 89.9% F1-score on EdNet, and 89.8% accuracy and 89.0% F1-Score on the OULAD dataset. This enhancement is obtained by integrating KSE-based concept-level knowledge estimation with ELRA-based educational optimization, which enables decisions to be modified to the learner's current knowledge and needs. Unlike existing approaches, which generally focus on a single capacity, such as tutoring, prediction, adaptation, multi-agent processing, or optimization, LITE-Trust collectively models learner state and optimizes personalized educational configurations. The reported results are expressed as mean ± standard deviation to indicate the performance variability across individual runs. For knowledge gain evaluation, the learner interactions in each

learning episode are divided into temporally ordered pre-learning and post-learning windows relative to the personalized learning process. Specifically, the pre-learning window comprised the assessment interactions preceding the recommended learning process, whereas the post-learning window comprised the assessment interactions following it. Knowledge gain (KG) is calculated from the enhancement in mean assessment performance between these two windows. These results are gained from KSE-driven recommendations and ELRA-guided optimized agent coordination, which minimize redundant educational actions and enhance learner-specific resource selection.

Table IX explains the performance of trustworthiness on EdNet and OULAD datasets. The LITE-Trust system achieves 94.5%/93.8% trust, 3.1%/3.4% hallucination, and 96.4%/95.6% factual consistency, demonstrating superior performance compared with the existing approaches. This achievement was obtained by integrating student profile

construction, KSE, ELRA, personalized agent coordination, and trust verification, which preserves learner-specific, context-guided, and factually interpretable decisions. Existing

approaches primarily focus on individual capabilities without integrating learner profiles with continuous trust evaluation, which results in lower trust and higher hallucination.

TABLE IX. ANALYSIS OF TRUSTWORTHINESS ON EdNet AND OULAD DATASET

Model	EdNet			OULAD		
	Trust (%)	Hallucination (%)	FC (%)	Trust (%)	Hallucination (%)	FC (%)
Conventional ITS	82.1±0.7	7.8±0.5	86.3±0.7	81.1±0.8	8.0±0.6	85.0±0.8
Conventional multi-agent system	82.9±0.8	7.5±0.5	86.7±0.7	81.8±0.8	7.8±0.5	86.1±0.7
LLM-based tutoring	84.5±0.8	7.0±0.5	88.2±0.6	83.6±0.9	7.3±0.5	86.9±0.7
Non-optimized multi-agent system	85.5±0.7	6.8±0.5	87.5±0.6	84.5±0.8	7.1±0.5	87.2±0.7
Adaptive Learning model	86.4±0.7	6.1±0.4	89.5±0.6	85.0±0.7	6.5±0.6	88.3±0.6
Optimization-based approach	87.1±0.6	5.9±0.4	89.8±0.5	86.1±0.7	6.3±0.6	89.1±0.6
LITE-Trust	94.5±0.4	3.1±0.3	96.4±0.4	93.8±0.5	3.4±0.3	95.6±0.5

TABLE X. ANALYSIS OF EXPLAINABILITY ON EdNet AND OULAD DATASET

Model	EdNet			OULAD		
	Fidelity (%)	Quality (%)	Interpretability (%)	Fidelity (%)	Quality (%)	Interpretability (%)
Conventional multi-agent system	80.5±0.8	79.3±0.9	76.9±0.7	79.8±0.7	78.5±0.7	76.2±0.8
Conventional ITS	82.3±0.7	80.6±0.8	78.7±1.0	81.6±0.9	79.8±0.8	77.8±0.8
Non-optimized multi-agent system	84.3±0.8	83.0±0.8	81.2±0.9	83.7±0.9	82.0±0.9	80.4±0.7
LLM-based tutoring	85.3±0.6	83.1±0.8	81.9±0.8	84.2±0.8	82.4±0.7	80.9±0.7
Optimization-based approach	86.5±0.7	84.5±0.8	83.4±0.8	85.1±0.8	84.2±0.9	82.3±0.8
Adaptive Learning model	86.9±0.7	85.6±0.7	84.1±0.8	85.8±0.8	85.0±0.8	83.5±0.6
LITE-Trust	95.0±0.4	94.1±0.5	92.6±0.5	94.6±0.5	93.4±0.5	92.1±0.5

Table X explains the explainability performance on EdNet and OULAD datasets. The LITE-Trust model achieves the highest fidelity of 95.0%/94.6%, explanation quality of 94.1%/93.4%, and interpretability of 92.6%/92.1% across two datasets. The highest results were achieved by implementing context-aware generation and validation of explanations, where the system combines learner-specific evidence and decision

outputs before generating the final explanations. Subsequently, the produced explanations remain close to the decisions actually made and are detailed enough to be informative for the learners. In contrast, existing methods often generate generalized explanations with restricted alignment to the learner-specific personalized decisions, which leads to lower fidelity, quality of explanation, and interpretability.

TABLE XI. ANALYSIS OF COMPUTATIONAL EFFICIENCY ON EdNet AND OULAD DATASET

Model	EdNet			OULAD		
	Latency (ms)	Memory (MB)	Parameters (M)	Latency (ms)	Memory (MB)	Parameters (M)
Adaptive Learning model	16.3	92.5	2.19	17.4	93.4	2.19
Conventional ITS	17.4	95.4	2.71	18.2	98.2	2.71
Non-optimized multi-agent system	18.4	98.7	3.08	19.5	99.7	3.08
Optimization-based approach	19.4	109.5	3.75	19.8	110.1	3.75
Conventional multi-agent system	19.5	105.8	3.43	20.4	106.9	3.43
LLM-based tutoring	20.2	118.1	6.20	21.5	119.8	6.20
LITE-Trust	10.5	80.5	1.92	9.5	81.3	1.92

Table XI depicts the computational efficiency of proposed and existing approaches on the EdNet and OULAD datasets in terms of latency, memory usage, and parameter count. LITE-Trust achieves the lowest latency of 10.5 ms and 9.5 ms, memory usage of 80.5 MB and 81.3 MB, and parameter count

of 1.92 M on EdNet and OULAD, respectively. Compared with existing methodologies, the proposed LITE-Trust system minimizes computational needs while preserving a lightweight architecture. The lowest latency demonstrates faster execution, while the reduced number of memory and parameters proves

the method's lightness. In addition, the similar results for different datasets signify that the efficiency of the LITE-Trust model is not tied to a particular dataset.

D. Earning Personalization Performance

Table XII illustrates the learning personalization performance of presented and existing works in terms of learner adaptation, knowledge-gap identification, recommendation accuracy, and learning-path quality. LITE-Trust system achieves the highest scores of 94.1%, 93.2%, 94.0%, and 92.9%, respectively. This enhancement is achieved by the combined effect of the presented approach, where the student profile construction extracts individual learner characteristics, KSE preserves the evolving knowledge state,

ELRA improves learning configurations, and responsive multi-agent architecture allocates learner-specific instructional tasks. The explainability engine then evaluates and transfers the rationale behind the produced learning decisions, while the trust-based strategy makes recommendations for trustworthy educational information. Unlike existing approaches, which generally deploy less integrated personalization strategies and generate restricted coordination among learner profiling, knowledge evaluation, optimization, recommendation, and decision validation. Therefore, they are less efficient in adjusting to changing learner needs, identifying specific knowledge gaps, and constructing high-quality personalized learning paths.

TABLE XII. ANALYSIS OF LEARNING PERSONALIZATION PERFORMANCE

Model	Learner Adaptation (%)	Knowledge Gap Identification (%)	Recommendation Accuracy (%)	Learning Path Quality (%)
LLM-based tutoring	81.2±0.7	79.9±0.8	80.5±0.7	78.7±0.8
Conventional multi-agent system	83.6±0.6	81.6±0.7	82.7±0.6	80.6±0.7
Conventional ITS	84.5±0.7	82.6±0.7	83.9±0.7	81.7±0.9
Non-optimized multi-agent system	85.3±0.6	83.4±0.6	84.2±0.6	82.6±0.7
Adaptive Learning model	86.4±0.5	84.2±0.7	85.1±0.7	83.4±0.6
Optimization-based approach	87.2±0.5	85.7±0.6	86.7±0.6	84.7±0.6
LITE-Trust	94.1±0.4	93.2±0.5	94.0±0.4	92.9±0.5

E. Trust and Hallucination Analysis

Table XIII depicts the trust and hallucination analysis of the proposed and existing approaches. The LITE-Trust model achieves the highest trust of 94.15% and hallucination reduction of 54.6%, which demonstrates the enhanced reliability of the produced educational outcomes. The overall trust score is calculated as the average of the dataset-specific

trust scores obtained on EdNet and OULAD. The hallucination reduction values are calculated using Eq. (33), with the LLM-based tutoring approach as the baseline. The results show a gradual improvement in hallucination reduction from the LLM-based tutoring method to the conventional multi-agent, conventional ITS, non-optimized multi-agent, adaptive learning, and optimization-based approaches, while LITE-Trust achieves the highest reduction.

TABLE XIII. TRUST AND HALLUCINATION ANALYSIS

Model	Trust Score (%)	Hallucination Reduction (%)
LLM-based tutoring	80.8±0.9	18.6±0.8
Conventional multi-agent system	82.5±0.8	22.5±0.8
Conventional ITS	83.7±0.7	25.3±0.7
Non-optimized multi-agent system	84.9±0.6	29.3±0.8
Adaptive Learning model	85.8±0.6	34.8±0.7
Optimization-based approach	86.8±0.5	39.8±0.6
LITE-Trust	94.15±0.4	54.6±0.5

The enhancement is qualified by the incorporation of student profile construction, KSE, ELRA, personalized multi-agent coordination, explainability engine, and trust-based validation approaches. These subsystems jointly support learner-specific decision generation, evidence-based recommendations, explanation evaluation, and controlled response generation. In contrast, the existing methods generally provide relatively restricted strategies for collectively evaluating generated content, learner-specific decisions, and

explanatory consistency. The higher hallucination reduction achieved by LITE-Trust therefore indicates improved control of potentially unsupported responses, contributing to higher overall trust in the generated educational outcomes.

F. Explainability Analysis

Table XIV presents the explainability performance of the proposed and existing approaches in terms of quality and transparency. Table XV represents the examples of

explanations and recommendations. The LITE-Trust model achieves superior performance of 94.1% explanation quality and 93.8% transparency, which shows that it provides informative and transparent explanations to the learner. The quality of the explanations and the transparency of the recommendations in the explanations were assessed using an automated scoring protocol. The explanations were scored and compared based on relevance, adequacy, clarity, consistency, and correspondence with the recommendation. For each explanation, scores were calculated based on these criteria and converted into percentages for comparison between the recommendation models. Additionally, the qualitative examples demonstrate that LITE-Trust is able to link the learner context to the recommended learning resources along with presenting an explanation for the recommendation outputs. This is made possible due to the combination of comprehensive learner profiling, knowledge-state representation, ELRA-based decision optimization, adaptive multi-agent coordination, and explainability engine mechanisms, which together facilitate exposing the recommendation factors. On the other hand, the compared

solutions provide only the recommendation output with little information about the factors influencing the outcome, thus being less transparent and informative in terms of explanation.

To ensure objective evaluation, fidelity, explanation quality, and interpretability are calculated using predefined evaluation criteria applied consistently across all generated responses. Fidelity measures the extent to which the explanation accurately reflects the evidence and reasoning associated with the generated response. Explanation quality evaluates relevance, completeness, clarity, and consistency of the explanation, while interpretability assesses whether the explanation can be readily understood and traced to the corresponding learning content and response. The evaluation criteria and scoring procedure are kept fixed across all experiments to avoid subjective variation. The reported percentages represent normalized scores aggregated across the evaluated responses. No post-hoc adjustment of scores is performed based on the obtained results. Where human evaluation is included, responses are independently assessed using the same predefined criteria by the designated evaluators, and the final score is obtained from their assessments.

TABLE XIV. EXPLAINABILITY ANALYSIS

Model	Explanation Quality (%)	Transparency (%)
LLM-based tutoring	80.5±0.7	79.6±0.7
Conventional multi-agent system	82.0±0.6	81.4±0.7
Conventional ITS	84.5±0.7	82.8±0.8
Non-optimized multi-agent system	85.6±0.6	84.7±0.7
Adaptive Learning model	86.9±0.5	85.4±0.6
Optimization-based approach	87.5±0.5	86.1±0.6
LITE-Trust	94.1±0.4	93.8±0.5

TABLE XV. EXAMPLE OF RECOMMENDATION AND EXPLANATIONS

Student Context	Recommendation	Generated Explanation
Low performance in algebra; good performance in arithmetic, fair performance in equations	Intermediate algebra learning resource	The learner is adequately proficient in arithmetic, but lacks knowledge in solving equations. Hence, it is recommended to study intermediate algebra to improve equation-solving skills
Poor understanding of probability, moderate performance in statistics fundamentals	Probability learning resource-intermediate level	The learner lacks an adequate understanding of the idea of probability. The recommended module was chosen based on the learner's current knowledge and the identified learning needs.

G. Lightweight Performance Analysis

Table XVI depicts the existing non-optimized multi-agent method with the full-agent, selective multi-agent, and overall LITE-Trust configurations. The evaluation in Table XVI is conducted on the OULAD dataset for the learner outcome prediction task. Full-agent LITE-Trust has more overhead since all agents are activated, while selective activation only triggers relevant agents according to educational requirements. The parameter counts presented in Table XVI relate only to the trainable parameters in the proposed LITE-Trust components, while the external GPT-4o mini model remains constant across all configurations of LITE-Trust. The difference between 2.46M and 3.35M in the full-agent configuration stems from

the activation of all task-specific agents and their corresponding processing layers. In the selective multi-agent configuration, only the agents relevant to the specified requirements are engaged, thus accounting for a lower count of trainable parameters at 2.46 M. The overall LITE-Trust configuration further decreases the parameter count to 1.92 M due to parameter sharing and the elimination of task-specific components within the integrated pipeline. Therefore, the variation in parameter counts originates from the proposed configurations rather than the external large language model. The 94.8% accuracy reported for the overall LITE-Trust configuration corresponds to the OULAD-specific result in Table VIII, whereas the 95.6% accuracy in Table VIII represents the separate EdNet evaluation.

TABLE XVI. LIGHTWEIGHT PERFORMANCE ANALYSIS

Method	Parameters (M) ↓	Latency (ms) ↓	Memory (MB) ↓	Accuracy (%) ↑
Non-optimized multi-agent system	3.08	18.4	98.7	88.7
Full multi-agent LITE-Trust	3.35	16.9	102.1	95.4
Selective multi-agent LITE-Trust	2.46	11.7	89.5	95.2
Overall LITE-Trust	1.92	9.5	81.3	94.8

The reported 1.92 M parameters represent only the locally implemented components of the proposed framework and exclude the externally accessed GPT-4o mini and embedding model. Therefore, the parameter count should not be interpreted as the complete deployed-system size. To provide a more realistic assessment of lightweight deployment, the evaluation considers end-to-end factors including local computational resources, API response latency, network communication overhead, external inference cost, and memory usage. The framework is therefore characterized as lightweight primarily with respect to its locally deployed components and selective agent execution, rather than as a parameter-light representation of the entire external AI service stack.

The overall LITE-Trust reduces computational costs by several means, such as compact learner and knowledge-state representations, limited ELRA search, selective agent activation, and efficient algorithms for processing explanations and trust. Selective agent activation contributes to efficiency by reducing unnecessary agent execution; however, the overall end-to-end enhancement cannot be attributed solely to reducing the number of active agents. Instead, the results indicate the agent-level reduction works together with the optimization of the complete pipeline, such as KSE, ELRA, retrieval, trust, and XAI processing. Therefore, selective activation provides an important agent-level efficiency benefit, while the overall lightweight performance results from the combined optimization of the entire LITE-Trust pipeline.

H. ELRA Optimization Performance

To ensure consistency in comparison, GA, PSO, ACO, LRA, and ELRA have been evaluated using the same population size, equal to 30, the same learner-specific fitness function, and the same experimental scenario with five independent runs using random seeds 11, 22, 33, 44, and 55. The value of 30 is selected for population size on the basis of the performed sensitivity analysis to balance the solution quality and computational effort. All the configurations used in the experiments can be considered reasonable as they conform to standard educational performance measures and use the

specific search and update mechanisms of the respective algorithms.

Table XVII depicts the comparison of the proposed and conventional optimization algorithms, namely GA, PSO, ACO, and LRA, in terms of convergence, fitness, standard deviation, computational time, agent activation, and learning performance. ELRA achieves faster convergence, higher fitness, the lowest standard deviation, lower computational time, fewer activated agents, and the best learning performance among the compared optimizers. The lowest standard deviation means high consistency of the results of independent runs. These advantages were achieved due to ELRA's learner-state-guided initialisation, constrained candidate generation and relevance-based optimization, and quality-and diversity-based selection, which aims to facilitate optimization by focusing the search on promising learner-specific configurations and avoiding redundant exploration. The overall comparison among five optimizers was statistically significant based on the Friedman test ($p \approx 0.001$), which indicates significant differences in optimizer performance across the evaluated methods.

The direct comparison with conventional LRA further demonstrates the contribution of the ELRA modifications. Under the same population size, learner-specific fitness function, experimental scenario, and independent random seeds, ELRA improves the best fitness from 0.89 obtained by LRA to 0.93, corresponding to a relative improvement of approximately 4.5%. ELRA also reduces convergence iterations from 38 to 28 (26.3% reduction), standard deviation from 0.025 to 0.012 (52.0% reduction), computational time from 15.4 ms to 10.8 ms (29.9% reduction), and average activated agents from 3.9 to 2.7 (30.8% reduction). Learning performance increases from 91.2% to 94.8%, an absolute improvement of 3.6 percentage points. These results indicate that the learner-state conditioning, trust-aware candidate filtering, and relevance-based search mechanisms provide measurable benefits beyond the conventional LRA search strategy.

TABLE XVII. ELRA OPTIMIZATION PERFORMANCE

Optimizer	Convergence iterations ↓	Best Fitness ↑	Standard Deviations ↓	Computational Time (ms) ↓	Activated Agents ↓	Learning Performance (%) ↑
GA	50	0.86	0.035	21.5	4.8	89.5
PSO	47	0.87	0.031	18.9	4.6	89.6
ACO	44	0.88	0.029	17.8	4.2	90.5
LRA	38	0.89	0.025	15.4	3.9	91.2
ELRA	28	0.93	0.012	10.8	2.7	94.8

I. Ablation Study

Table XVIII depicts the ablation study results, which highlight the contributions made by each component of the proposed LITE-Trust. Excluding ELRA reduces the learning pathway quality from 94.8% to 89.5% since the learner's KSE is not transformed into an optimized learning configuration, resulting in lower resource and path utility. This confirms ELRA's role in converting learner-state information into personalized educational configurations. Excluding the trust engine lowers trust by 9.8% because the produced recommendations are no longer validated against the reliable educational information repository via the trust-aware validation mechanism. This degradation reflects the role of trust evaluation in improving the reliability of educational decisions. Removing the explainability engine decreases the explanation quality from 94.6% to 78.3%, representing a 16.3% reduction, while slightly reducing latency from 9.5 ms to 8.8 ms. This trade-off is expected because removing XAI eliminates explanation-processing overhead but reduces transparency and interpretability. Excluding dynamic agent

selection increases the latency from 9.5 ms to 16.5 ms, representing a 73.7% increase, because unwanted agents are activated instead of directly employing only the relevant agents. These results demonstrate that LITE-Trust balances multiple objectives rather than maximizing every metric simultaneously, with XAI improving interpretability and agent selection enhancing computational efficiency. Finally, the overall performance of the complete LITE-Trust achieves 94.8% accuracy, 95.2% trust score, 94.6% XAI quality, 9.5 ms latency, and 94.8% learning path quality, since the student profile construction accurately represents the learners' features. KSE maintains the required knowledge state, ELRA determines the best configuration, dynamic coordination activates only necessary agents, and the trust and XAI engine enhance the transparency of decisions and recommendations. Therefore, the observed degradation after removing each component is consistent with its intended function, which confirms the complementary contribution of ELRA, trust evaluation, XAI, and dynamic agent selection to personalized, reliable, transparent, and efficient learning.

TABLE XVIII. ABLATION STUDY OF LITE-TRUST SYSTEM

Model	Accuracy (%)	Trust Score (%)	XAI Quality (%)	Latency (ms)	Learning-path Quality (%)
Model A: without ELRA	91.7	92.5	91.6	14.3	89.5
Model B: Without Trust Engine	93.8	85.4	91.9	9.7	93.0
Model C: Without XAI	94.0	94.2	78.3	8.8	93.6
Model D: Without Agent Selection	94.2	94.0	93.6	16.5	94.1
Model E: Overall LITE-Trust	94.8	95.2	94.6	9.5	94.8

To further clarify the contribution of individual modules, the ablation analysis is organized as controlled component-wise comparisons. Each variant removes one module while keeping the dataset, learner profiles, experimental conditions, and remaining framework components unchanged; changes in accuracy, trust, explainability, latency, and learning-path quality are compared with the complete model. The resulting LITE-Trust framework. This analysis identifies the principal contribution of each component while minimizing the influence of changes in other experimental factors. The observed improvements are interpreted as component-level contributions; interaction effects among multiple modules are considered a limitation of the current ablation design and are not claimed as independently established causal effects.

J. Sensitivity and Convergence Analysis

Table XIX shows that increasing the population size from 10 to 30 raises the fitness from 0.90 to 0.93 and the learning performance from 92.2% to 94.8%, while further increasing it to 40 provides only a marginal improvement to 94.9%. Similarly, increasing the number of iterations beyond 0 provides only marginal improvement in learning performance. Therefore, the selected configuration represents a performance-efficiency trade-off, where a population size of 30 and 60 iterations provides nearly the same performance as larger configurations while avoiding additional computational cost. Thus, the balanced configuration was optimized for both educational effectiveness and computational performance,

which reached the maximum fitness of 0.93 and a learning performance of 94.8%. These results suggest that ELRA is fairly robust to parameter variations, as moderate changes to the population size and iteration number affected the outcomes within a reasonable range. More importantly, population size, iteration number, and by extension, the optimization process itself are predefined computational constraints, not learner-specific variables. Hence, while ELRA extends the lightweight optimization principles of LRA, its educational adaptation introduces a restricted set of control parameters to control optimization and agent activation, while retaining the lightweight optimization philosophy of the original LRA. The weighting strategy is assessed using sensitivity analysis. Table XX below compares the trust, efficiency, learning, and balanced configuration based on attained fitness and learning performance. It can be seen that the balanced configuration had the highest fitness score of 0.93 and the best learning performance of 94.8% and was thus selected as the final ELRA configuration.

To examine whether the selected threshold values are robust, a sensitivity analysis is performed by varying the trust threshold and agent activation threshold across multiple levels while keeping the remaining experimental settings unchanged. The evaluation considers hallucination rate, response quality, agent utilization, latency, and learner performance. The same datasets, learner profiles, prompts, agent pool, and evaluation procedure are maintained for each threshold setting so that the

observed differences can be attributed primarily to threshold variation.

TABLE XIX. SENSITIVITY AND CONVERGENCE ANALYSIS

Parameter	Setting	Convergence iterations ↓	Fitness ↑	Learning performance (%) ↑
Population size	10	33	0.90	92.2
	20	29	0.92	93.7
	30	27	0.93	94.8
	40	27	0.93	94.9
Iterations	20	34	0.89	92.5
	40	29	0.91	94.1
	60	27	0.93	94.8
	80	26	0.92	94.9

TABLE XX. SENSITIVITY ANALYSIS OF ELRA WEIGHTING CONFIGURATIONS

Weighting Configuration	Trust	Efficiency	Learning	Fitness ↑	Learning Performance (%) ↑
Trust-based	0.49	0.20	0.30	0.88	92.5
Efficiency-based	0.20	0.49	0.30	0.89	92.8
Learning-based	0.20	0.29	0.49	0.90	93.4
Balanced	0.33	0.33	0.34	0.93	94.8

TABLE XXI. STATISTICAL JUSTIFICATION OF THE ELRA AGAINST THE GA, PSO, ACO, AND LRA

Trust threshold	Hallucination rate (%)	Response quality (%)	Agent utilization (%)	Latency (ms) ↓	Learner performance (%)
0.50	7.20	91.40	92.50	182	92.10
0.60	6.10	92.80	89.70	190	93.00
0.70	4.90	94.20	86.30	199	94.10
0.80	3.80	95.60	82.10	211	95.00
0.90	2.90	96.70	76.80	228	95.40

The threshold sensitivity results demonstrate the effect of threshold selection on the reliability and efficiency of the proposed framework. Lower thresholds allow more candidate responses and agents to participate, which can increase agent utilization and potentially reduce computational filtering, but may also permit less reliable responses to proceed. In contrast, higher thresholds impose stricter trust and activation requirements, which can reduce unnecessary agent execution and unsupported responses but may also reject useful responses or increase processing overhead associated with additional validation. The intermediate threshold setting provides a balanced operating point when considering hallucination control, response quality, agent utilization, latency, and learner performance.

K. Statistical Significance Analysis

Table XXI depicts the statistical justification of the ELRA against the GA, PSO, ACO, and LRA. The statistical unit is an

independent experimental run, with five runs conducted using seeds 11,22,33,44,55. Learner interactions and independent optimization evaluation are not treated as independent samples. The Friedman test reveals that there are significant differences between the five optimizers used ($\chi^2 \approx 18.65$, $p \approx 0.001$). The exact two-sided Wilcoxon signed-rank test is applied to the five paired run-level results, yielding $p \approx 0.063$ for the pairwise comparisons. The Nemenyi post-hoc analysis is based on the same five-run rank matrix. The effect sizes vary from 0.64 to 0.83, suggesting that the improvements are substantial. In addition, the 95% confidence interval 0.91-0.93 reflects the stability of ELRA fitness across the five runs. Thus, the statistical analysis uses run-level statistical units and avoids treating correlated learner or interaction observations as independent samples.

TABLE XXII. ANALYSIS OF STATISTICAL SIGNIFICANCE

Metrics	ELRA Evaluation	Result
Statistical unit	Independent experimental runs	5 runs (seeds 11,22,33,44,55)
Friedman test	Overall difference among GA, PSO, ACO, LRA, and ELRA	$\chi^2(4) \approx 18.65$, $p \approx 0.001$
Wilcoxon signed-rank	ELRA vs. GA/PSO/ACO/LRA	$p \approx 0.063$
Nemenyi post-hoc	ELRA vs. GA/PSO/ACO/LRA	$p > 0.05$
Effect size	ELRA vs. GA/PSO/ACO/LRA	0.83/0.79/0.71/0.64
95% Confidence Interval	ELRA Fitness	0.91-0.93

The statistical results have been clarified to avoid overstating significance (Table XXII). Although ELRA shows better mean performance than the compared optimizers, the pairwise Wilcoxon result of approximately 0.063 and the Nemenyi results above 0.05 do not demonstrate statistically significant superiority at the conventional 0.05 level. Therefore, ELRA is described as showing better observed performance and practical improvement, rather than statistically significant superiority in every pairwise comparison. The overall statistical analysis is interpreted together with effect sizes, confidence intervals, variability across independent runs, and mean performance to provide a more complete assessment of the observed differences.

L. Discussion

The results indicate that the improvements of LITE-Trust come from the complementary interaction of its proposed components, as further enabled by the ablation study; removing any individual module consistently degraded the corresponding performance metrics. ELRA enhances agent selection by leveraging the learner profile and KSE to determine the optimal learning configuration and identify relevant agents, thus restricting the optimization search to the learner-relevant educational configurations and avoiding unwanted exploration of unsuitable learning paths. Moreover, KSE identifies concept-level knowledge gaps and prerequisite mastery, which enables ELRA to allocate learning resources and educational agents based on the learner's actual learning requirements instead of depending on aggregate performance indicators. Trust Evaluation boosts performance by ensuring the factual

consistency of the results through confirming AI-generated educational content before delivery, thus minimizing the inclusion of unsubstantiated information and enhancing the reliability of educational recommendations. Dynamic orchestration facilitates reduced computational costs by limiting the engagement to necessary agents for each learning task, thereby avoiding unwanted agent invocations, which reduces inter-agent communication complexity and lowers processing time and memory consumption. This offers a favourable balance between learning performance and computational demands, as LITE-Trust performs better than all the baseline methods at a lower cost than multi-agent and optimization-based methods. The ablation results further indicate that the combined operation of KSE, ELRA, trust evaluation, explainability, and selective multi-agent orchestration generates more benefits than employing these components individually, which highlights their complementary contributions to personalization, optimization, trustworthiness, explainability, and computational efficiency. In comparison to the existing educational AI methods, which mostly focus on personalization, multi-agent orchestration, optimization techniques, explainability, or trust, LITE-Trust provides an integrated learner-centric approach that combines personalization, orchestration of multiple agents, optimization, explainability, and trust evaluation. In practice, such reduced complexity and selectivity would be critical to support the implementation in contexts where computational resources are limited, but learning personalization, performance, explanation, and trust are needed.

M. Limitations and Future Directions

Although LITE-Trust has shown promising results, there are still several limitations in terms of the current research. Firstly, it should be noted that it is based on a few educational benchmarks, making its ability to generalize questionable. Moreover, there was not enough practical application in a classroom setting, and using these lightweight GenAI components may lead to massive computational overhead. Then, it is hard to quantify the amount of learning gained by students since short-term gains do not always reflect long-term ones. Lastly, automatic trust and hallucination scores may miss many of the nuanced errors in natural language instruction.

The future directions of the research will address these concerns and explore different aspects of educational AI. In particular, they will focus on multimodal learning through speech, imagery, and behavioural data, and instruction via speech and vision. Federated educational AI and edge models will also be used to achieve better privacy and lighter computational needs while building more autonomous agents for education. Furthermore, future research will focus on longitudinal student modeling to understand learning gains better.

V. CONCLUSION

This study addressed the limitation of existing educational AI systems to provide personalized, optimized, trustworthy, explainable, and computationally efficient learning support. The LITE-Trust system integrated student profile construction, KSE, ELRA, dynamic multi-agent orchestration, explainability engine, and trust evaluation into a learner-centric pipeline.

First, ELRA enhanced learner-specific educational configurations by leveraging the learner profile and KSE to focus the optimization search on relevant learning paths, resources, and agent configurations while avoiding superfluous optimization steps and agent exploration. Second, the trust-aware and explainable mechanisms improved the measured trust and explanation quality of produced educational outcomes by validating educational information and providing learner-specific rationales. Third, the lightweight multi-agent orchestration mechanism exploited relevance-based selective agent invocations to prevent superfluous agent executions, thereby reducing the measured computational overhead and response latency.

In addition, the behaviour-based academic integrity analyzer complemented the framework by estimating indicators related to academic integrity and AI usage based on learners' interactions, thus providing educators with valuable information without actually measuring these variables directly. Overall, the experimental results demonstrated that the LITE-Trust system achieved higher measured personalization, trust, explainability, and computational efficiency over the compared baseline methods. The ELRA analysis revealed faster convergence, higher fitness, stable optimization, and fewer agent invocations, while the ablation and statistical tests highlighted the complementarity of the integrated components. Thus, the designed approach demonstrated the effectiveness of a combined learner-state modeling, educational optimization, selective multi-agent orchestration, explainability, and trust validation pipeline for the estimated educational learning scenarios. However, the evaluation is limited to the EdNet and OULAD benchmark datasets, and trust assessment is performed through an automated framework-based mechanism rather than long-term human validation. Furthermore, due to the absence of ground-truth labels in the data concerning academic integrity and reliance on AI, the results of the academic integrity analyzer should be regarded as behavioural estimates rather than measurements. Large-scale real-world classroom deployment and comprehensive learner satisfaction measurement were not conducted. These aspects remain important directions for future evaluation.

FUNDING SOURCES

This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia, Grant No.: KFU265137.

ACKNOWLEDGMENT

The authors are thankful to the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, and Al-Fayha Private College, Jubail, Saudi Arabia, for providing financial assistance.

DATA AVAILABILITY

The datasets used and/or analyzed during the current study are available in the Kaggle repository, <https://www.kaggle.com/datasets/anlgrbz/student-demographics-online-education-dataoulad> and

<https://www.kaggle.com/datasets/anhtu96/ednet-kt34>. The datasets were obtained from their respective publicly accessible sources and used in accordance with their applicable terms of use and data-access policies. The datasets support the learner interaction, assessment, activity, and knowledge-estimation analyses reported in this study. The data used for the experiments are available from the corresponding public dataset repositories.

DECLARATIONS

Ethics Approval and Consent to Participate: This study used publicly available educational learning analytics datasets, namely the Open University Learning Analytics Dataset (OULAD) and EdNet. No participants were directly recruited or contacted by the authors, and no new personally identifiable information was collected for this study. The analyses were conducted on the datasets as provided by their respective sources and in accordance with the applicable terms of use and data-access requirements. Therefore, no additional participant recruitment or consent procedures were undertaken by the authors.

REFERENCES

- [1] Reddig, J.M., Arora, A., & MacLellan, C.J. (2025). Generating In-Context, Personalized Feedback for Intelligent Tutors with Large Language Models. *International Journal of Artificial Intelligence in Education*, 35, 3459 - 3500. <https://doi.org/10.1007/s40593-025-00505-6>
- [2] Ahmed, Z., Shanto, S.S., Rime, M.H., Morol, M.K., Fahad, N., Hossen, M., & Abdullah-Al-Jubair, M. (2024). The Generative AI Landscape in Education: Mapping the Terrain of Opportunities, Challenges, and Student Perception. *IEEE Access*, 12, 147023-147050. <https://doi.org/10.1109/ACCESS.2024.3461874>
- [3] Dahal, C., Chen, J., Ye, M., & Xiong, Z. (2026). Multi-Agent Systems in Education: A Survey from the Trustworthiness Perspective. *IEEE Transactions on Artificial Intelligence*. <https://doi.org/10.36227/techrxiv.176704820.01980102/v1>
- [4] Youssef Alyoussef, I., Mohammed Drwish, A., Adel Albakheet, F., Hassan Alhajhoj, R., & Ahmed Al-Mousa, A. (2025). AI Adoption for Collaboration: Factors Influencing Inclusive Learning Adoption in Higher Education. *IEEE Access*, 13, 81690-81713. <https://doi.org/10.1109/ACCESS.2025.3567656>
- [5] Algerafi, M.A., Zhou, Y., Alfadda, H., & Wijaya, T.T. (2023). Understanding the Factors Influencing Higher Education Students' Intention to Adopt Artificial Intelligence-Based Robots. *IEEE Access*, 11, 99752-99764. <https://doi.org/10.1109/ACCESS.2023.3314499>
- [6] Oh, S. (2025). Integration of MATH41 and Generative AI in Pre-Service Mathematics Teacher Education: An Empirical Study on Lesson Design Competency. *IEEE Access*, 13, 128959-128973. <https://doi.org/10.1109/ACCESS.2025.3586593>
- [7] Pradeesh, N., Thushara, M.G., Arun Krishna, K., Pranav, V., & Krishnamoorthy, S. (2025). AI-Driven Personalized Learning and Remedial Recommendation Through Knowledge Concept-Centric Evaluation. *IEEE Access*, 13, 207817-207837. <https://doi.org/10.1109/ACCESS.2025.3638427>
- [8] Ma, Y., Tang, X., & Huang, X. (2025). AI-Powered Adaptive English Language Learning Systems: Leveraging Deep Learning Algorithms and Natural Language Processing for Personalized Teaching Approaches. *IEEE Access*, 13, 153189-153198. <https://doi.org/10.1109/ACCESS.2025.3603602>
- [9] Lou, H., Yue, P., & Chen, J. (2025). Research on Adaptive Education Path Dynamic Programming Algorithm Based on Reinforcement Learning and Cognitive Graphs. *IEEE Access*, 13, 105446-105462. <https://doi.org/10.1109/ACCESS.2025.3578389>
- [10] Viswanathan, N., Meacham, S., & Adedoyin, F.F. (2022). Enhancement of online education system by using a multi-agent approach. *Comput. Educ. Artif. Intell.*, 3, 100057. <https://doi.org/10.1016/j.caeai.2022.100057>
- [11] Chang, C., Shih, S., & Kuo, B. (2026). A Dual-Agent Intelligent Tutoring System for Teaching Two-Step Fraction Word Problems. *IEEE Transactions on Learning Technologies*, 19, 636-647. <https://doi.org/10.1109/TLT.2026.3708780>
- [12] Rida, Z., Boukachour, H., Ennaji, M., & Machkour, M. (2025). Dynamic Adaptation in an Intelligent Multi-Tutoring System: A Multi-Agent Approach. *IEEE Access*, 13, 124587-124601. <https://doi.org/10.1109/ACCESS.2025.3576651>
- [13] Evangelista, E., & Salman Bukhari, S.M. (2026). From ethical principles to executable governance: A policy-as-code framework for trustworthy AI in higher education. *Computers and Education: Artificial Intelligence*. <https://doi.org/10.1016/j.caeai.2026.100594>
- [14] Raji, N.R., Kumar, R.M., & Biji, C.L. (2024). Explainable Machine Learning Prediction for the Academic Performance of Deaf Scholars. *IEEE Access*, 12, 23595-23612. <https://doi.org/10.1109/ACCESS.2024.3363634>
- [15] Lu, Y., Wang, D., Chen, P., & Zhang, Z. (2024). Design and Evaluation of Trustworthy Knowledge Tracing Model for an Intelligent Tutoring System. *IEEE Transactions on Learning Technologies*, 17, 1701-1716. <https://doi.org/10.1109/TLT.2024.3403135>
- [16] Radharani, S., Sistla, V., & Kishore Kolli, V.K. (2026). CSBAM-MobileNet: a lightweight attention-enhanced deep learning model for student attentiveness classification. *Engineering Science and Technology, an International Journal*. <https://doi.org/10.1016/j.jestech.2026.102326>
- [17] Yu, X., & Tian, Y. (2025). Enhancing Computer Education through IoT-Enabled Learning Environments Leveraging Mobile Edge Computing for Real-Time Feedback. *Systems and Soft Computing*. <https://doi.org/10.1016/j.sasc.2025.200433>
- [18] Qorich, M., & Ouazzani, R.E. (2026). Context-based automatic question generation for educational assessment via lightweight transformer augmentation. *Eng. Appl. Artif. Intell.*, 177, 114953. <https://doi.org/10.1016/j.engappai.2026.114953>
- [19] Kalwar, A., Shekhawat, D., Joshi, S., & Bairwa, A.K. (2025). A Genetic Algorithm Approach to Adaptive Learning: Enhancing Learner Satisfaction and Performance. *IEEE Access*, 13, 102815-102829. <https://doi.org/10.1109/ACCESS.2025.3578867>
- [20] Yuhana, U.L., Djunaidy, A. and Purnomo, M.H., (2024). Development of dynamic personalized learning paths based on knowledge preferences and the ant colony algorithm. *IEEE Access*, 12, pp.144193-144207. <https://doi.org/10.1109/ACCESS.2024.3442312>
- [21] Zeng, Y., & Xu, X. (2022). The Construction and Optimization of an AI Education Evaluation Indicator Based on Intelligent Algorithms. *Int. J. Cogn. Informatics Nat. Intell.*, 16, 1-22. <https://doi.org/10.4018/IJICINI.315275>
- [22] Prasad G, Arun, et al. "Artificial Intelligence in Computer Science: An Overview of Current Trends and Future Directions." *Advances in Artificial and Human Intelligence in the Modern Era*, edited by S. Suman Rajest, et al., IGI Global Scientific Publishing, 2023, pp. 43-60. <https://doi.org/10.4018/979-8-3693-1301-5.ch002>
- [23] Hesham Mohammed Ali Abdullah, Elham Alzain. AI-driven Predictive Analytics for Early Identification of At-risk Students Academically in Saudi Universities., *International Journal of Intelligent Engineering and Systems*, Vol.19, No.6, 2026 DOI: 10.22266/ijies2026.0630.05
- [24] Mosleh, Mohammed A. S. , et al. Adaptive Mirage Search Optimization for Educational Workload Scheduling in Multi-Cloud Environments. *International Journal of Advanced Computer Science & Applications*, 2026, v. 17, n. 7, p. 240