Explainable AI for Hyper-Personalized Learning: Personalized Intelligent Tutoring Systems

Shivani Sharma[1] and O. P. Rishi[2]

[1*, 2] Department of Computer Science & Informatics, University of Kota, Kota, INDIA
[1*]`shivanisharma@uok.ac.in`
[2]`dr.oprishi@uok.ac.in`

Abstract.
Artificial Intelligence (AI) plays a vital role in facilitating an Intelligent Tutoring System. XAI-based personalised Intelligent Tutoring Systems (ITS) for lifelong learning are transforming education by offering adaptive and customised learning experiences. But the black-box nature of many AI learning systems makes it difficult to understand how decisions are made. This dearth of transparency diminishes user trust and raises concerns about the ethical and moral justification of using such systems. Particularly sensitive student information and high-stakes suggestions are of concern. This paper proposes a multi-modal framework that integrates ITS. The approach analyzes multiple learner signals, including keystroke patterns and facial expression signals, to better understand learner behavior and provide interpretable explanations for the system's decision.[1] The study describes the existing gaps that constitute a layered system architecture, models cases in field scenarios, addresses privacy and prejudice, and defines a roadmap for deploying previously workable ITS solutions.

Keywords: XAI, Hyper-personalized learning, Intelligent Tutoring Systems (ITS), Learning analytics, Interpretable machine learning, Shapley, LIME, Transparency, Trust, Governance.

1 Introduction

Intelligent Tutoring Systems powered by Artificial Intelligence (AI). Lifelong, personalized education (ITS) is transforming education. Nevertheless, this is restricted by their black-box nature and ethical considerations, particularly where sensitive student data and high-stakes recommendations are involved. In [6], the author proposes a multi-modal approach that integrates with an Intelligent Tutoring System (ITS) to enhance the learning process. Keystroke expression Explainable Artificial Intelligence (XAI) using keystroke signals, facial signals, logs of expressions, emotions, and behavior to propel transparent, hyper-personalized scale learning. The research details the current gaps, offers a layered architecture of systems that replicate case scenarios, addresses bias and privacy concerns, and proposes a wayfinder to effectively and ethically deploy EdTech. solutions. In some cases, both teachers and students may find it difficult to understand the logic behind the system's recommendations of a particular lesson plan or revision plan. Because of this lack of justification, the administrators may find it difficult to determine how to ensure fairness or regulatory compliance in audits, enable

B. Singh et al. (eds.), *Proceedings of the International Conference on Advances in Computing Technology and Artificial Intelligence (COMPUTATIA 2026)*, Atlantis Highlights in Intelligent Systems 18,
https://doi.org/10.2991/978-94-6239-713-2_53

teachers to intervene effectively, and hold students responsible for their own development. Researchers and practitioners have turned to studying it more often. XAI is an important component in helping human management with individualized learning upon the discovery of this gap [6]. Explainable Artificial Intelligence (XAI) can improve transparency in AI-based systems, helping to build user trust and facilitate better collaboration between intelligent tutoring systems and human instructors [1], [9], [11].

AI-powered solutions by means of offering fair and easy-to-understand explanations, e.g., pointing out the patterns of work by a student, which resulted in a recommendation. In the context that follows, we combine modern XAI approaches informed by diverse learners' data sources. Besides live adaptation to students' needs, the system was designed to answer the why and how behind recommendations. Our research shows three significant attributes: a scalable software architecture framework that connects different learner data sources to XAI interfaces; a mechanism for aggregating inter-readable explanations that are human- and person-readable; and evidence that such transparency is enhanced. It offers better results that increase teacher acceptance more useful information. Balancing transparency and accuracy is essential for XAI-based ITS, while deploying AI-driven learning platforms is essential for supporting effective and reliable teamwork in high-stakes learning environments [6].

2 Related Work

2.1 Opacity in Educational AI

Teachers, students, and academic leaders expect AI systems to provide fair and clear advice. When AI suggests classroom transformations, such as changes in instruction delivery, extra support for learning, or assessment adjustments. Then, the classroom transformation should explain the story of every decision taken in the scenario on which it is based. But many adaptive and intelligent systems rely on complex patterns of teaching and learning from past performance or behaviour. These systems never explain how they work. This kind of problem, due to a lack of clarity, makes it hard for teachers, instructors, and educators to understand student learning behaviour and the learning process. Teachers, instructors, educators, and education policymakers fret that these systems often focus on surface-level participation and might reinforce biases. The reinforcement biases limit the reasons behind their suggestions and miss deeper learning needs. Also, without transparency, teachers struggle to determine whether recommendations align with the curriculum and students' interests, and make informed choices. As a result, trust in ITS and AI tools for teaching & learning could weaken over time. Hence, the concerns about fairness, especially in decisions about student progress or course eligibility, add to this loss of confidence [3, 4].

Recent studies indicate that unintentional biases can affect the data selected and, in turn, how machine learning models are built [3, 4]. Even well-defined systems can exhibit these biases due to the selection of data and models developed from it [3, 4]. The reason for giving such importance is not only to create accurate AI-based intelligent tutoring systems but also to ensure they clearly explain and justify their decisions to everyone involved. Researchers, developers, and stakeholders need to work together to prioritize transparency, accountability, and fairness throughout the development of these systems.

2.2 Taxonomies of Interpretability

Researchers who are working on machine learning often consider two main aspects to make models easier to understand. The first aspect uses models that are naturally interpretable because their structure is straightforward. For example, these models include simple decision trees and rule-based systems during decision-making. The advantage of selecting such naturally interpretable models is that the system users can clearly see how different input features influence the final result or recommendation. However, focusing too much on interpretability can sometimes lower the model's accuracy, especially with complex or highly non-linear data [3,4].

The second aspect is post hoc interpretability, which seeks to explain how complex "black-box" models make decisions after training. This training involves using specialised techniques to interpret the model's prediction mechanism. SHAP and LIME XAI-based personalised models both use specialised techniques to interpret the model's prediction mechanism. SHAP is developed based on ideas from cooperative game theory to assign importance scores to individual features, while LIME explains specific predictions by creating simpler local versions of the original model. In AI-based education and ITS, where advanced models can uncover hidden patterns in students' learning, these explanation methods can be very helpful. But their usefulness depends on whether educators and students can review the explanations and trust the reasoning behind the system's recommendations.

2.3 Model -Agnostic Techniques in XAI-Based ITS

Techniques used for personalised learning, such as LIME and SHAP, have become very important in educational technology. These personalised learning techniques offer a variety of methods of explanation. When a user uses LIME to understand their learning, LIME makes predictions after a successful learning session. The LIME has one key benefit: it can explain individual predictions one at a time. Also, it can show why a system suggested a specific problem or provided a student with a hint by creating a simpler, easier-to-understand model near that example. This kind of clarity helps teachers to spot unusual system actions and give useful feedback when students have questions about their personalised assignments.

As SHAP is developed based on the principles of cooperative game theory to assign importance scores to individual features, it operates in parallel by equitably allocating the "credit" or influence of each input feature to both individual decisions and broader trends in learning behaviour and individual learning. Teachers who want to know which behavioral factors are driving key recommendations for an entire class or curriculum unit have found the visual summaries useful. Both approaches offer significant benefits, according to recent studies, and real-world classroom experiments have demonstrated the effectiveness of data-driven and explainable approaches in educational systems [5] [7][8]. SHAP permits more comprehensive audits of system logic and potential biases in content delivery, while LIME facilitates focused, learner-specific explanations.

3 Proposed System Architecture

The hype around personalized learning is driven by the use of XAI. The framework is made up of five interconnected layers. provide transparency, flexibility, and scalability in a variety of. educational contexts.

3.1 Multi - Model data Integration Layer

The bottom layer records heterogeneous learner signals. through multiple channels. Behavioral Analytics monitors the behavioral trends, such as mouse movements and scroll velocity. click sequences and dwell times of given content elements. Keystroke Dynamics refers to typing records, correction rates, and pauses, which are assessed as indicators of cognitive load and will be used to deduce cognitive load engagement levels. The Affective Computing Components' facial expressions are analyzed using convolutional neural networks (CNNs) to detect learners' emotional states. The model is trained on facial expression datasets and evaluated using standard validation techniques to ensure reliable emotion recognition. [11]. The system uses privacy-preserving methods by processing data on edge devices that store sensitive biometric data and sending it. only derived features, not audio or visual streams. Environmental Data Integration is used in Contextual Data Integration. time of day, session duration, device type, and location to enhance personalization context. The multidimensional methodology enables analysis of multiple learner states, including cognitive performance and behavioral indicators, thereby providing a more comprehensive understanding of learner behavior [6][7][8]. learner states and needs.

3.2 Dynamic Learner Modelling Engine

The basic modeling system tracks the entire set of learner profiles, which continuously evolve in response to multi-mode inputs. Cognitive State Estimation uses assessment performance together with behavioral measures to estimate the learner's current knowledge and learning state within the dynamic learner modelling engine [6][8]. Learning Style Adaptation determines the visual, auditory, or kinesthetic content-delivery preference through analysis of interaction patterns. Changes in the learning trajectory are modeled in Temporal Modeling, taking into account effects such as fatigue, variations in motivation, and environmental factors related to time of day or session length. Bayesian updating mechanisms are used to update the system with a set of uncertainty estimates, enabling more powerful personalization decisions. Peer interaction is taken into account in Social Learning Integration, patterns, and collaborative learning actions to improve the accuracy of individual modeling and to provide a social context. recommendations.

4 Proposed Methodology

The proposed methodology integrates high-quality personalization, clear, shareholderoriented communication, and management checks. It includes three stages that interact with and are mutually constructive to develop an all-inclusive XAI framework, as depicted in Fig. 1.

4.1 Stage 1: Black Box Personalization and Multi- Model Profiling

This phase builds adaptive policies on a multi-modal basis. Formation of learning evidence and definition of the basis of the following explanation generation.

- **Data Induction and Processing:** Learning Traces. (responses, attempts, time-on-task, hint usage), navigation behavior (scroll/Click patterns) and optional affective displays (e.g., frustration inferred from interaction abnormalities) are collected with overt consent and in accordance with confidentiality constraints. Data cleaning, normalization, and aggregation are performed at the session and unit levels. The focus of feature selection is on interpretability and generalizability.[3]
- **Adaptive Policy Training:** A high-performing model, such as gradient-boosted trees or a neural network–based policy model, is trained using the collected learner interaction data. The model parameters are optimized during training, and its performance is evaluated using standard validation techniques to assess its ability to handle sequencing big data, difficulty scoring, hint timing, and modality adaptation [7][8].
- **Synthetic Rules to Discovery of Founder Behavior:** The trained model analyzes comprehensive interaction logs using a policy-based approach to discover patterns and generate interpretable rules and a coded corpus describing both the features of the learner state and the needed actions [7] [8]. Interpretable surrogates can be used to extract human-readable rules using this corpus [3][5].
- **Multi-Modal Feature Engineering:** Complex feature. The processing of raw interaction data into interpretable representations using extraction techniques. Keystroke dynamics are transformed to measure cognitive load, mouse motions to engagement and patterns of attention, and temporal sequences. trajectories.

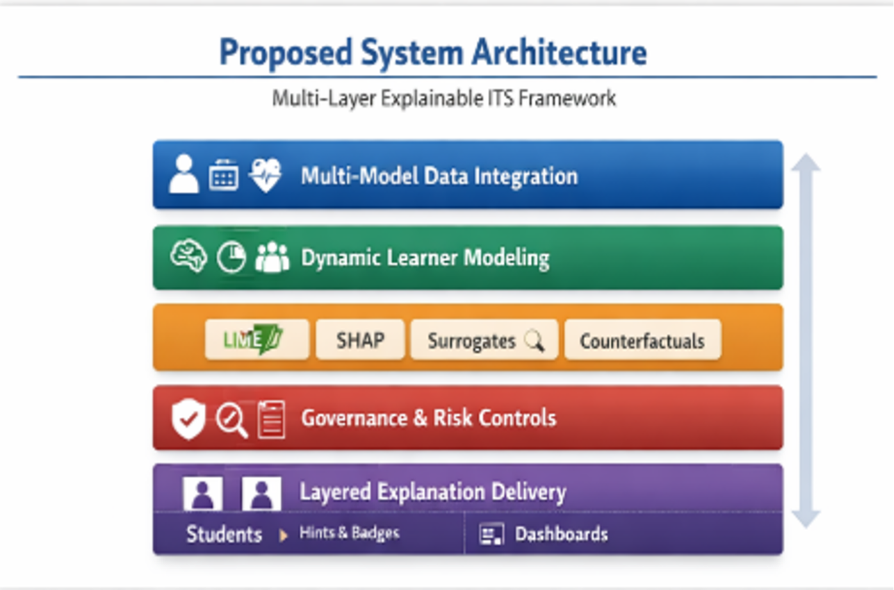

Fig 1. ITS Framework

4.2 Stage 2: Post-Hoc Interact Expounding and Stakeholder-Centered Account

This phase brings about local and worldwide transparency. Trading performance: Multiple explanation modalities are designed to support the diverse needs of stakeholders, such as students, instructors, and system developers, by providing interpretable insights into the model's decisions [9][11].

- **Local Explanations via LIME:** This is the one where a set of recommendations (such as assign Concept A before Concept B, and so on) can be explicitly explained. surrogate issues which feature (recent errors, prolonged). dwell, repeated hint requests) exerted the greatest influence on the decision, presented in the form of a brief logical or graphical explanation.
- **SHAP global attributions: Cohort-level SHAP summary plots monadically reveal determinants of factor adaptation**. (e.g., place is mastered by a prior, time- on task-frequency control time pacing). These insights help designers identify incompatible proxies and re-optimize.
- **Interpretable Surrogate Modeling: Decision trees or lists of rules, which are trained on the behavioral corpus.** The compact and auditable rules give approximations of black-box decisions. These artifacts, written in human-legible formats, aid in governance review and course design iteration.
- **Layered Explanation Delivery:** Scale explanations down to brief words and badges on behalf of learners (Recommended this module because of recent

mastery holes on recursion), and scale upwards to more elaborated dashboards on instructors (feature attributions, administrator level audit trails, confidence, counterfactuals), and confidence (bias checks, drift indicators).

- **Generations of Counterfactuals:** Generating explanations. Invent what-ifs to help learners demonstrate their variation. Different recommendations may be provided based on behaviors or responses, which promote metacognitive awareness and self-regulation.

4.3 Stage 3: Governance, Consistency, and Risk Controls

This phase discusses the rollouts, which should be accompanied by far-reaching measures to ensure protection and accountability and to ensure effective and ethical operation.

- **Consistency Checks of Policies:** Scanning Surrogate rules for differences, diffuse demarcations, or within-group inequity. Draftsmen classification observes changes in feature distributions and term-delinquent policy behavior. Privacy-by-design: Sensitive modalities (webcam/audio) are locally determined, optional, and processed; systems store processed features rather than unprocessed media. Users can inspect and revoke data modalities.

- **Bias and Fairness Audits:** The distribution of Attribution is used to compare subgroups to identify disparate influence patterns; counterfactual probes test whether small, irrelevant influence patterns are present in a patient; recommendations change with changes.
- **Continuous Monitoring and Alerting:** Real-time monitoring solutions track the quality of explanations, model performance, fairness metrics, and other indicators, and warn administrators about potential problems that require intervention.

5 Experimental Setup monitoring solutions monitor the quality of the explanations, model performance, etc., and fairness metrics, and

5.1 Data and Task

The analysis of course modules that applied adaptive learning was conducted, including scheduling and one-on-one suggestions in subjects such as science, mathematics, and programming. Included information. It was conducted at three universities, anonymized over a 12-week semester. Standard tests and interaction logs were taken. Students were provided with feedback generated from both affective signals (such as emotional states) and analytical insights derived from their interaction data, enabling a more personalized learning experience; timing proxies also fit well with the models [6],[10]. The experiments were conducted both in the laboratory and in the classroom, which verifies real-world relevance. The dataset consisted of interaction data collected from 450 students participating in the learning activities. The data were collected via the Intelligent Tutoring System during classroom sessions, and participants were selected from the study sample to analyze learner behavior and system performance [5], [7].

5.2 Models and Tooling

To achieve personalization in the baseline models, high-capacity models, such as combinations of boosted trees and deep networks, were evaluated and selected for their performance and robustness. LIME and SHAP. It was calculated on demand with optimized, scalable libraries. Surrogate interpretable models, such as decision trees and rule lists, are used to consider policies back-end) enabled students and instructors to check their information, suggestions, and individual commentaries. Data governance had FERPA and institutional practices. no tendentious gibe or video. Hybrid cloud/edge on the computation end only: a hybrid cloud/edge solution guarantees privacy and performance. Sensitive inferences were created on devices; training and analytics used resource-intensive, safe cloud clusters and general-purpose processors.[2]

5.3 Evaluation Metrics

The system was evaluated using quantitative performance metrics and user experience measures, including learning performance indicators and user feedback on the system's usability and effectiveness [5], [10]. Learning Effectiveness: Pre- and post-assessments, mastery rate, and 30-day retention were used to gauge learning in problem-solving.

- User interfaces: Subjective lucidity of explanations, in general, student trust rating, acceptance of recommendations, and so on. Cognitive load when receiving feedback was observed.
- Stakeholder Metrics: The administrators analyzed compliance and bias-mitigation accuracy; instructors concentrated on improvement. Student-monitored metacognition, timing, and insight of intervention and self- control development. The study focused on course modules that used adaptive pacing and personalized hinting in subjects like science, math, and programming. The study included information on students who participated during a semester at three local universities. During data capture, anonymized interaction logs and standard tests were used. Instructors had the option to collect data in an effective, advanced analytical manner, but clickstream and timing proxies also performed well with the models. Students have to provide all the learning freedom while capturing data for the study.

Studies were conducted in the lab and classroom to verify real-world relevance. The total sample consisted of 450 students from graduate and undergraduate cohorts with diverse experimental backgrounds. All these are given in Fig.2

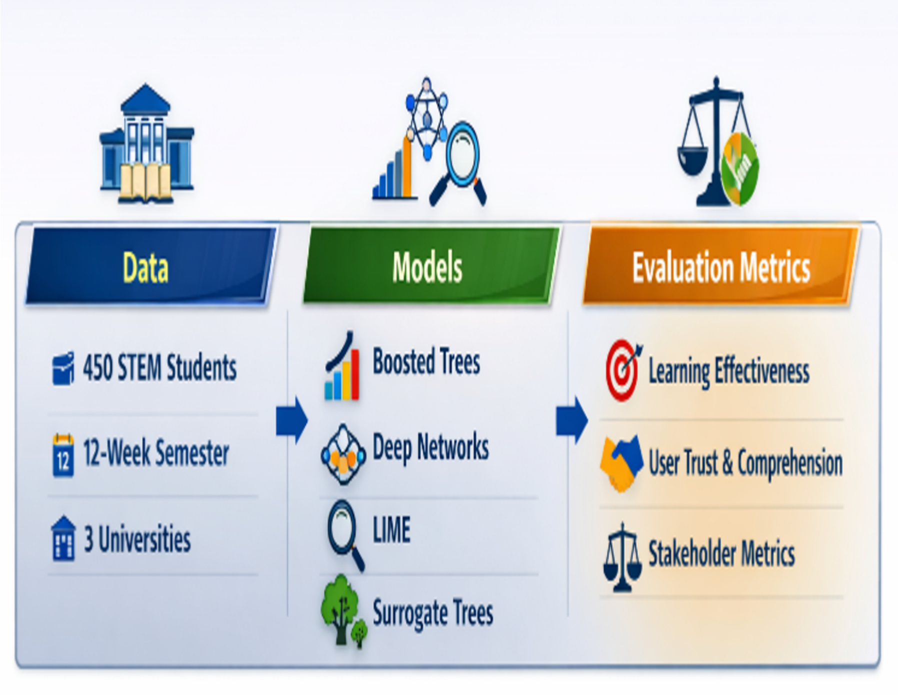

Fig.2. Experimental Setup.

6 Experimental Results

6.1 Usability and Comprehension

The results of the experiments indicated that XAI-enhanced feedback substantially improved comprehension and apparent trust. For example, the usability evaluation indicated that participants reported higher levels of trust and understanding after interacting with the system, with the average scores showing an increase of approximately 1.6 times based on user feedback collected through structured usability questionnaires, which is compared to the black-box condition, the evaluation detected several missed bias cases in the black-box mode that were disregarded under the XAI configuration. The results also show 3 missed bias cases in black-box mode and 0 in XAI mode [9][11].

6.2 Case Scenario

In the experimental case *a,* learner Shivani working on a programming problem is presented to demonstrate how the proposed system analyzes learner interactions and provides explanations. This example represents a typical scenario from the broader experimental dataset, and during learning, her behavior signals stress and confusion [6][7]. During the experiment, the system responded based on the students' behavior. For explanation, take one example - Shivani’s behavior according to facial signals and provide an explanation visualization. The explanation visualization states that "Shivani, you are struggling with the Programming Problem, and your facial expression and inactivity increased, which triggered additional scaffolded practice and visual aids". This type of feedback helped Shivani to connect his behavioral patterns to the system's recommendations and adjust his study strategies accordingly.

Table 1. Comparative Metrics - Comparison between Black Box ITS and XAI-ITS

Metric	Black Box ITS	XAI-ITS
Recommendation Acceptance	54%	79%
User Trust Score	2.7/5	4.3/5
Bias Cases Detected	1/10	4/4

In the black box ITS, 54% of students followed the system's suggestions, whereas in the XAI-based ITS scenario, 79% followed the recommendations. Similarly, Black Box has below-average trust, whereas XAI-based ITS has high trust (4.3 out of 5). The same scenario applies to Bias Case detection.

7 Result and Discussion

7.1 Quality of Personalization and Learning Effectiveness

There was an improvement in black-box policies. statistically significant differences in over static baselines of measures of learning. The decline in mastery was an average of (180.42), of which mathematics showed the most significant gains. problem-solving modules (reduction of a quarter). Pass rates on target concepts increased by 15 percent overall, with larger gains among students who initially had rankings below the median (a 22 percent increase). The system supported by the XAI ensured the quality of personalization, added transparency, and statistical gains in learning equivalent with black-box conditions only ($p = 0.05$, paired T-test). This finding resolves the issues raised regarding explanation generation and its detrimental effect on learning efficacy. Knowledge retention Analysis: The analysis conducted at a 30-day follow-up indicated a 12 percent increase in retention among students who used it. XAI-based systems, when compared with black-box counterparts, provided better explanations in transparent systems, which led to enhanced learning and the development of metacognition.

- **Faith, Belief, and Pedagogical Conscientiousness**

  Conditioning with XAI led to much higher recommendation acceptance (78 systems) and trust ratings among users (4.2/5.0 vs. 2.8/5.0). Students said they understood the reason for this much better. Next step," with 87. The teachers noted a better assessment of student challenges, and 92 percent of the educators stated that XAI dashboards improved their pedagogical insights. Early diagnosis of struggling. Students achieved 40 percent better results than traditional analytics, enabling earlier interventions. Explanation Comprehensibility: Student interviews demonstrated that brief, action-based explanations were preferred over elaborated technical explanations. Optimal explanation length in-context hint was 15-25 words, and in-context was 50-75 words, reflection prompts [2,10].

- **Global Pattern and Design Revision**

  SHAP analysis showed consistent global in-cohort drivers: patterns of recent errors (35 frequency 28), suggesting that the system has been able to adapt to current performance rather than past performance. The audit identified an issue with overreliance on engagement proxies, as some programming modules were extended rather than interpreted as misunderstandings, and debugging time was mistaken for productive struggle. This understanding led to the development of policy and enhanced the precision of recommendations. Temporal Pattern Discovery: SHAP temporal analysis identified regular patterns in learning effectiveness, and morning sessions showed greater vulnerability to challenging lessons that would require more support [2].

- **Privacy and Ethical Reflections**

  Recent developments in education via Explainable AI focus not only on transparency but also on ethical implications, including bias mitigation, user privacy, and user control. Algorithmic bias can unconsciously support inequities in society by disproportionately shaping recommendations for some demographic groups. Therefore, it is essential to introduce ethical audits and fairness checks into the XAI pipeline to achieve fair treatment across learners. Privacy-preserving algorithms, i.e., local processing of powerful fine-grain consent management and sensitive data users' personal information, create trust and adherence to laws such as FERPA and GDPR. Consent and Agency: Students used granular consent outputs to choose which modalities they could be collected from; full multi-modal analytics was advanced by 85 percent of participants who had seen the benefits of the explanation.[2] The data was inspected through data inspection tools. A third of the respondents mainly wanted to learn how their learning was going. Recommendations were affected by patterns.

- **Human-Computer Interaction: XAI Systems**

  The intelligent tutoring systems should be effectively explainable, taking into consideration cognitive load and user experience. Layered explanations, providing summaries on an overview of answers to students and in-depth educator pedagogic diagnostics, maximize utility and comprehension. Conversational and other interaction modalities. Leader and counterfactual reasoning assist in a more profound learner. metacognition, which will render AI recommendations not only transparent but also pedagogically useful. Development of explanatory interfaces aligned with education. workflows boost the effect of adoption and instruction [8].

- **Risk Controls Governance**

  Attribution parity tests detected changes in engagement proxies over time in high-stress conditions. (midspan, finals) allowing for automatic recalibration of personalization algorithms. Monitoring of drift revealed curriculum evolution in feature distributions, as suggested by retraining protocols. Layered explanations had a positive effect in lessening the cognitive burden, with simple descriptions requiring, on average, 8 seconds of study time, and detailed viewing of the instructor, on average, 45 seconds. Personalizing an explanation based on previously learned technical knowledge increased clarity by 23. Scalability Test: The system was tested to scale to between 10,000 parallel users under load, archetypal use cases, and explanation generation latency under a second, via sub-second average load duration, distributed computing, and optimized forms of caching [6,5].

8 Limitations and Future Work

There are a number of significant limitations of the current approach, which indicate new directions for studies.

Explanation Quality Dependencies: The quality of explanations in teaching and learning largely depends on effective feature engineering that extracts relevant system features and on accurate data representation. Noisy behavioural indicators during data extraction and representation can lead to misleading conclusions that confuse users. Therefore, future research should focus on developing explanation methods that remain reliable even when data quality varies.

Contextual Generalization: While the teaching-learning framework has been tested in various educational settings, further validation is needed to assess its performance across various educational scenarios, cultures, conditions, and institutions [9].

The educational framework for teaching and learning emphasises cross-cultural research. It ensures explanations remain relevant across diverse scholarly beliefs and traditions. This study's evolution focused on immediate learning outcomes and user experience. evaluations of other factors, such as career development and lifelong educational outcomes, require longitudinal research to capture the sustained impact of acquired skills. Large language models can be used for explanations. The capabilities of large language models may facilitate more natural explanations. Advancing causal explanation involves moving beyond correlational feature characteristics to elucidate how learning occurs. Here, the key considerations include the system life cycle, expanding the granularity of bias audits to enhance privacy protection measures, and increasing user control and transparency in artificial intelligence-based personalisation [9].

9 Conclusion

The educational stakeholders want to eliminate the gap between powerful personalisation algorithms and transparency. Explainable Artificial Intelligence (XAI) enables closing the gap between powerful personalisation algorithms and transparency. This framework supports productive, effective education while boosting trust, acceptance, and teaching insights across multiple instructional settings. We offer a clear method for adding XAI to personalised learning systems, back it up with tests showing that explanations work well among different groups, and deliver practical advice for protecting privacy. The present research shows that adding explanations does not reduce the quality of personalisation when done carefully. The advancement in AI in key areas, using layered explanations, focusing on stakeholders, and monitoring bias, are effective ways to deploy AI responsibly. Future educational XAI-based ITS should integrate transparency with performance to keep personalisation effective and to ensure accountability for learners and teachers. This research provides a foundation for achieving that goal while maintaining the sophisticated capabilities that make AI-driven personalization so powerful.

References

1. Ribeiro, M.T., Singh, S., Guestrin, C.: Why should I trust you? Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2016), pp. 1135–1144 (2016)
2. Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems (NeurIPS 2017), pp. 4765–4774 (2017)
3. Rudin, C.: Stop explaining black box models for high-stakes decisions and use interpretable models instead. Nat. Mach. Intell. 1(5), 206–215 (2019)
4. Holstein, K., Vaughan, J.W.: Improving fairness in machine learning systems: What do industry practitioners need? In: Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI 2019), pp. 1–16 (2019)
5. Papamitsiou, Z., Economides, A.A.: Learning analytics and educational data mining in practice: A systematic literature review of empirical evidence. Educ. Technol. Soc. 17(4), 49–64 (2014)
6. Graesser, A.C., Conley, K., Olney, A.: Intelligent tutoring systems. In: Handbook of Research on Learning and Instruction, pp. 451–473 (2012)
7. Baker, R.S., Inventado, P.S.: Educational data mining and learning analytics. In: Learning Analytics, pp. 61–75 (2014)
8. Koedinger, K.R., Stamper, J.C., McLaughlin, E.A., Nixon, T.: Using data-driven discovery of better student models to improve student learning. In: Proceedings of the 16th International Conference on Artificial Intelligence in Education (AIED 2013), pp. 421–430 (2013)
9. Conati, C., Porayska-Pomsta, A., Mavrikis, M.: AI in education needs interpretable machine learning: Lessons from open learner modelling. arXiv preprint arXiv:1807.00154 (2018)

10. Kochmar, E., Yankovsky, D., et al.: Automated data-driven generation of personalized pedagogical interventions in intelligent tutoring systems. Int. J. Artif. Intell. Educ. 32(2), 323–349 (2022)
11. Karpouzis, K.: Explainable AI techniques for intelligent tutoring systems (2024)