Adaptive Learning in Higher Education: A Knowledge Tracing and Explainable AI Approach

Nandan Reddy Muthangi, University of Toledo

Nandan Reddy Muthangi is a senior-year Computer Science and Engineering student at the University of Toledo, where he currently serves as the President of the Association for Computing Machinery (ACM) Student Chapter. He is actively involved in research and development across multiple interdisciplinary domains. Nandan works as a Research Assistant in the LONG and the Cyber-Physical Human Systems (CPHS) Lab, where his contributions span projects involving autonomous drone navigation, swarm-based multi-object tracking, and intelligent sensing systems.

In addition to his research efforts, he is part of a senior design team developing UniMatch, a smart college recommendation system that leverages Elasticsearch and KNN-based algorithms for personalized university suggestions based on SAT scores, academic interests, and extracurricular profiles.

This summer, Nandan will be expanding his research interests into the field of computational techniques in materials science, exploring how machine learning and simulation tools can be used to model, predict, and optimize material behavior at the nanoscale. He is on track to graduate in the current semester and will be joining Cornell University's Master of Engineering program in Materials Science and Engineering in Fall 2025.

Ananya Singh, The University of Toledo

Ananya Singh is a Bachelor's student at the University of Toledo, majoring in Computer Science and Engineering. She is an undergraduate research assistant at the RIM Lab, where her research focuses on machine learning and its applications in IoT. Her work includes integrating AI with IoT systems to develop innovative solutions for real-world problems such as wildfire detection, where she led the development of drone-based sensing systems and predictive analytics for early fire alerts.

Ananya has co-authored research paper in the areas of explainable AI, autonomous systems, and drone technology. She is also actively involved in student leadership as the Treasurer of the ACM Chapter and Technical Head of the ACM-W Chapter at the University of Toledo, organizing professional development events and tech workshops.

She will be joining Duke University in Fall 2025 to pursue her Master's in Electrical and Computer Engineering, with a focus on Machine Learning and Big Data. Her long-term goal is to bridge the gap between education and technology, fostering advancements in personalized and accessible learning platforms powered by artificial intelligence.

©American Society for Engineering Education, 2025

(WIP) Adaptive Learning in Higher Education:
A Knowledge Tracing and Explainable AI Approach

Abstract

The rapid advancements in Artificial Intelligence (AI) are transforming the educational landscape, offering new opportunities to enhance learning and teaching experiences. This work-in-progress research focuses on the application of Deep Knowledge Tracing (DKT) and Feedforward Neural Networks (FNN) to model and predict student performance effectively. The models were evaluated using two datasets, EdNet and Student Performance, to track and analyze students' learning progress. The DKT model demonstrated strong predictive performance on the EdNet dataset, achieving a test accuracy of 94.34%, while the FNN performed well on the Student Performance dataset within its tabular context. SHAP-based interpretability techniques were applied to validate the model's predictions, revealing critical factors influencing performance, such as prior grades and socio-economic variables.

While this study primarily focuses on the development and evaluation of predictive models, future work aims to integrate these models into an AI-powered educational platform. This envisioned system will provide personalized learning experiences for students and actionable insights for educators by leveraging explainable AI (XAI) techniques. By bridging the gap between predictive modeling and real-world educational applications, this research lays the groundwork for intelligent, adaptive, and transparent educational technologies.

Keywords: Deep Knowledge Tracing, Artificial Intelligence, Explainable AI(XAI).

1. Introduction and Related Work

Recent advancements in Artificial Intelligence are set to drive transformative changes across various domains, including healthcare [1], environmental sciences [2], business management [3], and most notably, education [4]. The keyways AI is being used in the field of education include personalized learning, Intelligent tutoring systems (ITS), optimizing administrative processes, and enhancing accessibility and engagement. By tailoring learning experiences to individual student needs, AI-powered systems have the potential to increase engagement, improve academic performance, and provide more equitable access to education.

Among the emerging AI-driven methodologies, Deep Knowledge Tracing (DKT) [9] has gathered attention as a robust approach for tracking and predicting a student's evolving knowledge state. DKT leverages sequential data, such as a student's interactions with learning materials, to predict future performance and enable adaptive learning experiences. However, while DKT offers strong predictive capabilities, its black-box nature and lack of interpretability limit its adoption in real-world educational systems.

While approaches like the three-layer knowledge tracing model proposed by Y. Lu et al. [11] aim to enhance trustworthiness in intelligent tutoring systems, they fall short in providing sufficient transparency in model predictions. Similarly, the work by Ma et al. [12] underscores that although prior studies have successfully integrated AI into Intelligent Tutoring Systems (ITS) and adaptive learning platforms, achieving a balance between high model performance and explainability remains a persistent challenge.

To address these limitations, recent efforts have explored the integration of Explainable AI (XAI) methods, particularly SHAP (SHapley Additive exPlanations) [8], which can uncover the most influential features driving a model's predictions. The combination of predictive accuracy and interpretability is thus crucial for building effective and trustworthy AI-powered learning systems.

This study explores the effectiveness of predictive modeling in two different educational data settings. We evaluate models on both a sequential interaction dataset (EdNet) and a static tabular dataset (Student Performance) to understand how different architectures, such as FNN and DKT, perform under varying data structures. Rather than comparing these models head-to-head, the goal is to assess each model's suitability within its appropriate context and to examine how explainability techniques like SHAP can enhance their interpretability.

2. Methodology

This research evaluates the performance of DKT and FNN on two datasets: EdNet and Student Performance, focusing on predictive accuracy and feature-level interpretability. By applying SHAP, we analyze the key determinants of student performance, such as prior grades and socio-economic variables, providing actionable insights for educators and researchers.

2.1. Datasets

This study utilizes two publicly available datasets: EdNet-KT1 and the Student Performance dataset, selected for their complementary characteristics. EdNet-KT1, part of the larger EdNet corpus collected from a Korean online learning platform, contains over 131 million student interactions. We sampled 200 session log files from KT1 for exploratory data analysis, revealing a mean session length of 23.4 and a median of 11, making it highly suitable for sequential modeling. KT1 logs question-answer interactions with associated timestamps, bundle identifiers, and elapsed time, enabling the use of models like Deep Knowledge Tracing (DKT) that rely on temporal dependencies. In contrast, the Student Performance dataset includes 649 students from Portuguese secondary schools, with structured features such as prior grades (G1, G2), parental education, study time, and school support. The final grade (G3) was binarized for classification into high (≥10) and low (<10) performers. As it lacks time-based logs, this dataset was used with non-sequential models such as Feedforward Neural Networks (FNN). Together, these datasets allow for the evaluation of predictive modeling in both sequential interaction and static tabular learning environments.

2.2 Data Preprocessing

For the EdNet dataset, preprocessing involved several steps to prepare the data for modeling. Initially, missing values were removed to ensure data consistency. Numerical transformations were applied using label encoding to convert categorical features, such as question_id and user_answer, into numerical representations. Additionally, a new binary column, correct, was created to indicate whether a student answered a question correctly or incorrectly. For the DKT model, student interactions were grouped by user ID, generating sequences of questions and correctness, which were then padded to uniform lengths for input into the model.

Preprocessing the Student Performance dataset focused on feature scaling and data preparation for classification tasks. Numerical features, including prior grades and attendance, were normalized to ensure consistent input. The dataset was split into training and testing sets using an 80-20 ratio to enable robust evaluation. In cases where class imbalance was observed, techniques such as oversampling or weighting the loss function were considered to ensure balanced learning.

2.3 Model Training

To model and predict student performance, two approaches were employed: FNN and DKT model.For the EdNet dataset, the FNN consisted of an embedding layer to map question_id to dense vectors, followed by dense layers with 128 and 64 units activated by ReLU functions. The output layer employed sigmoid activation to predict the binary correct label. The model was trained using the Adam optimizer with a binary cross-entropy loss function, a batch size of sixty-four, and over 10 epochs. A 20% validation split was used to monitor model performance, and overfitting was mitigated through early stopping based on validation loss.

The DKT model for the EdNet dataset leveraged an LSTM-based architecture to manage sequential data. The model began with an embedding layer for question_id, followed by an LSTM layer with 128 units to capture temporal dependencies in student interactions. Dropout (30%) and L2 regularization were applied to prevent overfitting. The model was trained over 50 epochs, with early stopping and a learning rate scheduler (ReduceLROnPlateau) employed to optimize training and halt when validation performance plateaued. These techniques ensured robust generalization and reduced the risk of overfitting.
The Student Performance dataset followed a similar modeling pipeline, with slight adaptations to manage its tabular nature. Key features such as G1, G2, and absences were fed into the models. For the DKT model, sequences of student data were padded to uniform lengths, and dropout layers were utilized to maintain model robustness. Both models were evaluated using the same hyperparameters as the EdNet dataset.

2.4 Evaluation Metrics

The models were evaluated using standard classification metrics, including accuracy, precision, recall, and F1-score, to ensure a comprehensive assessment of performance. Additionally, SHAP analysis was conducted to interpret the model predictions and identify key contributing features, such as the impact of prior grades and socio-economic factors on student performance. The

combination of high performance and interpretability underscores the practical applicability of these models in educational settings.
Metrics:

$$\text{Accuracy} = \frac{TP+TN}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\text{F1 Score} = 2 \times \frac{\Pr ecision \times Recall}{\Pr ecision + Recall}$$

3. Results and Discussion

We evaluated the performance of two datasets, EdNet and Student Performance, using two models: FNN, and DKT. The results focus on model accuracy, precision, recall, F1-score, and interpretability through SHAP analysis. Additionally, training and validation trends are presented to demonstrate model convergence and generalization capabilities.

3.1. Feedforward Neural Network on the EdNet Dataset

The FNN achieved a test accuracy of 81.94% and a test loss of 0.4118, demonstrating moderate predictive performance. The training and validation accuracy and loss graphs (Figure 1a, 1b) illustrate smooth convergence, with training accuracy reaching 82% and validation accuracy stabilizing at 81%, while both losses steadily decrease over 10 epochs. The classification report (Table 1) highlights a weighted F1-score of 0.82, indicating overall reliability in predictions, with particularly high precision for "Incorrect" responses. However, the relatively lower recall for "Correct" responses suggests some limitations, likely due to imbalances in the dataset or the simplicity of the mode.

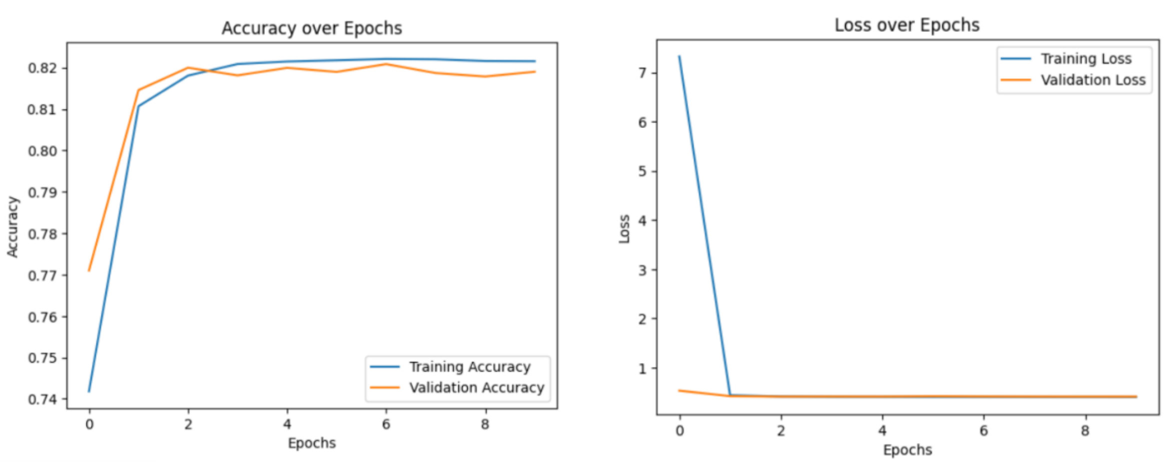

Figure 1. (a): Training and Validation Accuracy over Epochs (b) Training and Validation Loss over Epochs

Metric	Incorrect	Correct	Macro Avg	Weighted Avg
Precision	0.86	0.69	0.78	0.81
Recall	0.89	0.62	0.76	0.82
F1-Score	0.88	0.65	0.77	0.82

Table 1: Classification Metrics for Feedforward Neural Network on EdNet Dataset

3.2. Deep Knowledge Tracing (DKT) on EdNet Dataset

The DKT model demonstrated superior performance on the EdNet dataset, achieving a test accuracy of 94.34% and a test loss of 0.1373, significantly outperforming the FNN. The model's training and validation accuracy and loss graphs (Figure 2a, 2b) reveal rapid convergence, with validation accuracy stabilizing near the test accuracy, and validation loss remaining low throughout training. These trends indicate effective learning with minimal overfitting.
The classification report (Table 2) highlights the model's robust performance across both "Correct" and "Incorrect" categories, with balanced precision, recall, and F1-scores. The macro F1-score of 0.82 underscores the model's overall reliability, with notable improvements in recall for correctly predicting "Correct" responses compared to the FNN.

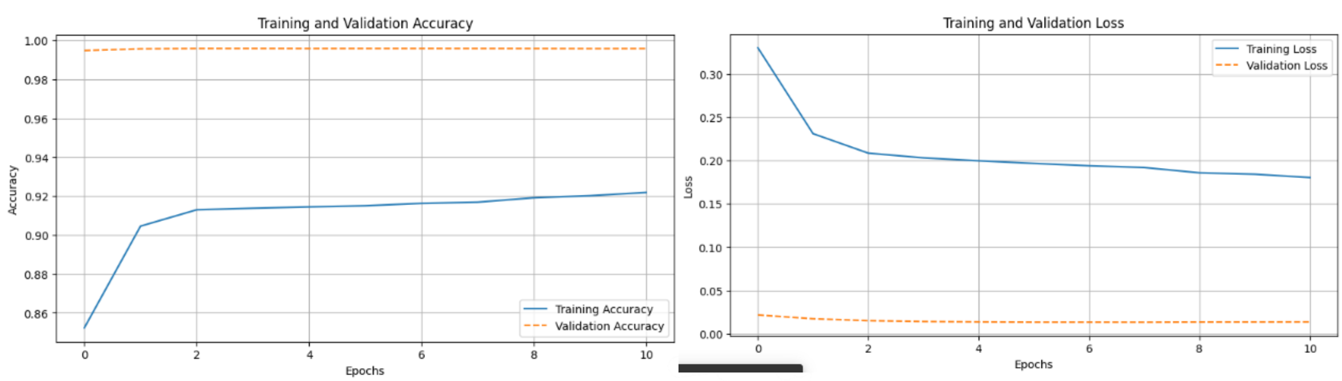

Figure 2. (a): Training and Validation Accuracy over Epochs (b) Training and Validation Loss over Epochs

Metric	Incorrect	Correct	Macro Avg	Weighted Avg
Precision	0.96	0.73	0.85	0.94
Recall	0.98	0.62	0.80	0.94
F1-Score	0.97	0.67	0.82	0.94

Table 2: Classification Metrics for Deep Knowledge Tracing (DKT) on EdNet Dataset

3.3. Feedforward Neural Network on Student Performance Dataset

The FNN performed well on the Student Performance dataset, achieving a training accuracy of 88.92%, validation accuracy of 91.92%, and test accuracy of 91.92%. As shown in Figures 3a and 3b, training and validation trends converged smoothly over 50 epochs, with minimal overfitting. Validation accuracy stabilized near the test accuracy, while losses decreased steadily, reflecting strong generalization capabilities.

The classification report (Table 3) highlights the model's reliability, with a precision of 97% and an F1-score of 95% for high-performing students, and a recall of 87% for low-performing students. Overall, the macro F1-score of 87% and weighted F1-score of 92% demonstrate its balanced performance. Predicted grades closely matched actual grades in regression tasks, with minor deviations (Table 4). SHAP analysis further validated the model by identifying prior grades (G1, G2) and parental education (Medu, Fedu) as key predictors, aligning with established educational research.

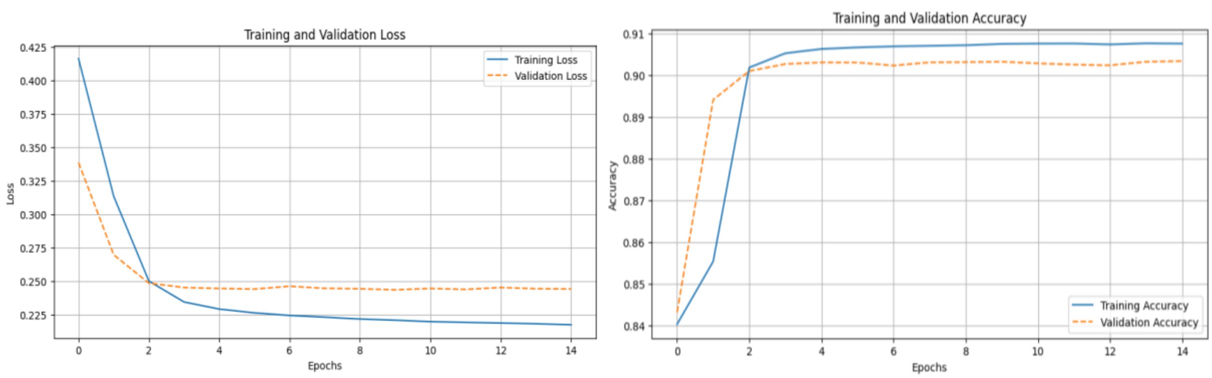

Figure 3. (a): Training and Validation Accuracy over Epochs (b) Training and Validation Loss over Epochs

Metric	Low	High	Macro Avg	Weighted Avg
Precision	0.73	0.97	0.85	0.93
Recall	0.87	0.93	0.90	0.92
F1-Score	0.79	0.95	0.87	0.92

Table 3: Classification Metrics for Feedforward Neural Network on Student Performance Dataset

Student	Actual Grade (Denormalized)	Predicted Grade (Denormalized)
1	11	9.92
2	15	14.11
3	8	7.21
4	10	10.94
5	10	10.54

Table 4: Predicted Vs Actual Grades

3.4. Deep Knowledge Tracing (DKT) on Student Performance Dataset

The DKT model demonstrated strong performance on the Student Performance dataset, achieving a test accuracy of 94.34% and a test loss of 0.1373, indicating effective optimization and minimal overfitting, with low and stable validation loss reflecting the model's generalization capabilities.

The classification metrics (Table 5) emphasize the model’s ability to differentiate between high- and low-performing students. For low-performing students, the model achieved a precision of 96%, recall of 98%, and an F1-score of 97%, ensuring reliable identification. Although performance for high-performing students was comparatively weaker, with an F1-score of 67%, the overall weighted F1-score of 94% highlights the robustness of the model. Feature importance analysis using SHAP (Figure 4) identified prior grades (G1 and G2) as the most significant predictors, followed by parental education (Medu, Fedu) and attendance (absences), aligning with established educational research on academic performance determinants.

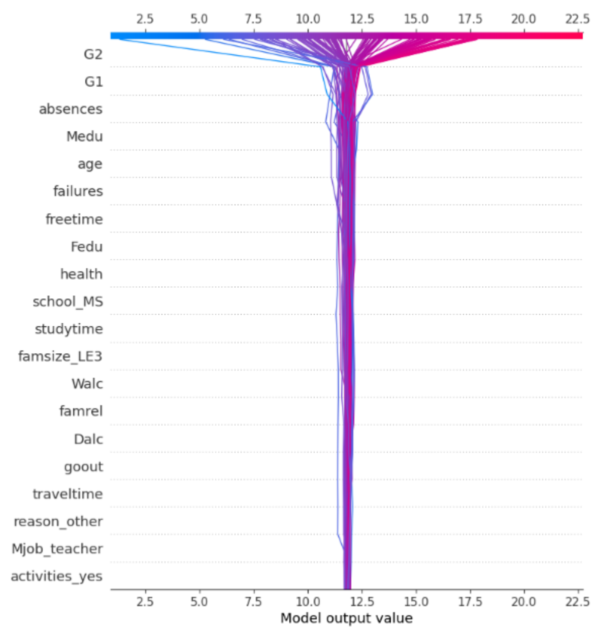

Figure 4. SHAP Dependence Plot Showing Key Feature Contributions to Model Output for Student Performance Dataset

Metric	Low	High	Macro Avg	Weighted Avg
Precision	0.96	0.73	0.85	0.94
Recall	0.98	0.62	0.80	0.94
F1-Score	0.97	0.67	0.82	0.94

Table 5: Classification Metrics for Deep Knowledge Tracing (DKT) on Student Performance Dataset

3.5 Discussion

The results demonstrate the effectiveness of both the FNN and DKT models in predicting student performance across two datasets, with DKT consistently outperforming FNN in terms of accuracy and generalization. On the EdNet dataset, DKT achieved a significantly higher test accuracy and better balance across performance metrics, underscoring its ability to model sequential data effectively. On the Student Performance dataset, both models delivered high overall accuracy, with FNN excelling in precision and recall for high-performing students, while DKT provided more balanced predictions for low-performing students. SHAP analysis further validated the interpretability of the models, highlighting the critical influence of prior grades and parental education. These findings suggest that leveraging advanced models like DKT, combined with explainability techniques, can offer valuable insights for educational interventions and personalized learning strategies, paving the way for data-driven improvements in academic outcomes.

4. Conclusion and Future Work

This study evaluated Deep Knowledge Tracing (DKT) and Feedforward Neural Networks (FNN) on two educational datasets, EdNet and Student Performance, to assess their predictive capabilities and interpretability through SHAP. By applying both models across sequential and tabular data types, we found that DKT performed best in modeling temporal learning behaviors, while FNN excelled at leveraging structured academic and demographic information. Rather than direct alternatives, the two models offer complementary strengths, and SHAP-based analysis revealed common influential features such as prior grades, attendance, and parental education.

These findings inform our broader vision of an AI-powered educational platform that integrates both behavioral and contextual modeling to deliver personalized learning support. Such a system would interface with existing Learning Management Systems (LMS) through APIs to provide real-time insights for both students and educators. Future work will focus on refining the DKT model, addressing class imbalance, and testing the platform across more diverse datasets.

Practical testing will be a key component, involving real-world classroom settings to evaluate usability, gather feedback, and ensure the system aligns with educational needs. Pilot studies with students and educators will guide interface refinement and assess the platform's effectiveness in fostering personalized learning. Additionally, ethical AI practices, including privacy safeguards and explainability, will remain a priority to ensure trust and fairness in decision-making. By bridging predictive modeling with real-world applications, this research lays the foundation for a transformative AI-driven educational ecosystem.

7. References

[1] S. Lopes, G. Rocha, and L. Guimarães-Pereira, "Artificial intelligence and its clinical application in Anesthesiology: a systematic review," *J. Clin. Monit. Comput.*, vol. 38, no. 2, pp. 247–259, 2024.

[2] A. Konya and P. Nematzadeh, "Recent applications of AI to environmental disciplines: A review," *Sci. Total Environ.*, vol. 906, p. 167705, 2024.

[3] R. Thakur, "Introduction to artificial intelligence and its importance in modern business management," in *Leveraging AI and Emotional Intelligence in Contemporary Business Organizations*, IGI Global, 2024, pp. 133–165.

[4] J. Mao, B. Chen, and J. C. Liu, "Generative artificial intelligence in education and its implications for assessment," *TechTrends*, vol. 68, no. 1, pp. 58–66, 2024.

[5] O. Onesi-Ozigaguan, Y. J. Ololade, N. L. Eyo-Udo, and D.O. Ogundipe, "Revolutionizing education through AI: A comprehensive review of enhancing learning experiences," *Int. J. Appl. Res. Soc. Sci.*, vol. 6, no. 4, Apr. 2024.

[6] P. Cortez, "Student Performance Dataset," *UCI Machine Learning Repository*, Nov. 26, 2014. [Online]. Available: https://archive.ics.uci.edu/dataset/320/student+performance

[7] D. Shin and S. Lee, "EdNet," *GitHub Repository*, 2019. [Online]. Available: https://github.com/riiid/ednet

[8] S. M. Lundberg, "An introduction to explainable AI with Shapley values," *SHAP Documentation*. [Online]. Available: https://shap.readthedocs.io/en/latest/example_notebooks/overviews/An%20introduction%20to%20explainable%20AI%20with%20Shapley%20values.html

[9] C. Piech *et al.*, "Deep knowledge tracing," in *Adv. Neural Inf. Process. Syst.*, vol. 28, 2015.

[10] A. M. Hasanein and A. E. E. Sobaih, "Drivers and consequences of ChatGPT use in higher education: Key stakeholder perspectives," *Eur. J. Investig. Health Psychol. Educ.*, vol. 13, no. 11, pp. 2599–2614, Nov. 2023, doi: 10.3390/ejihpe13110181.

[11] Y. Lu, D. Wang, P. Chen, and Z. Zhang, "Design and evaluation of trustworthy knowledge tracing model for intelligent tutoring system," *IEEE Trans. Learn. Technol.*, vol. 17, pp. 1661–1676, 2024, doi: 10.1109/TLT.2024.3403135.

[12] A. Marouf, R. Al-Dahdooh, M. J. Abu Ghali, A. O. Mahdi, B. S. Abunasser, and S. S. Abu-Naser, "Enhancing education with artificial intelligence: The role of intelligent tutoring systems," *Int. J. Eng. Inf. Syst. (IJEAIS)*, vol. 8, no. 8, pp. 10–16, 2024.